arXiv:1405.0740v2 [cs.DS] 5 Nov 2014

Hardness of Graph Pricing Through Generalized Max-Dicut

Euiwoong Lee*

Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213.

Abstract

The Graph Pricing problem is among the fundamental problems whose approximability is not well-understood. While there is a simple combinatorial $\frac{1}{4}$-approximation algorithm, the best hardness result remains at $\frac{1}{2}$ assuming the Unique Games Conjecture (UGC). We show that it is NP-hard to approximate within a factor better than $\frac{1}{4}$ under the UGC, so that the simple combinatorial algorithm might be the best possible. We also prove that for any $\epsilon > 0$, there exists $\delta > 0$ such that the integrality gap of n^{δ}-rounds of the Sherali-Adams hierarchy of linear programming for Graph Pricing is at most $\frac{1}{4} + \epsilon$.

This work is based on the effort to view the Graph Pricing problem as a Constraint Satisfaction Problem (CSP) simpler than the standard and complicated formulation. We propose the problem called Generalized Max-Dicut(T), which has a domain size $T+1$ for every $T \geqslant 1$. Generalized Max-Dicut(1) is well-known Max-Dicut. There is an approximation-preserving reduction from Generalized Max-Dicut on directed acyclic graphs (DAGs) to Graph Pricing, and both our results are achieved through this reduction. Besides its connection to Graph Pricing, the hardness of Generalized Max-Dicut is interesting in its own right since in most arity two CSPs studied in the literature, SDP-based algorithms perform better than LP-based or combinatorial algorithms — for this arity two CSP, a simple combinatorial algorithm does the best.

*Supported by a Samsung Fellowship, Venkatesan Guruswami's US-Israel BSF grant 2008293, and NSF CCF-1115525. Most of this work is done while visiting Microsoft Research New England. euiwoonl@cs.cmu.edu

1 Introduction

Consider the following natural problem for a seller with a profit-maximization objective. The seller has n types of items $1, \ldots, n$, each with unlimited copies, and there are m customers $1, \ldots, m$. Each customer j has her own budget b_j and a subset of items $e_j \subseteq \{1, \ldots, n\}$ that she is interested in. Customers are *single-minded* in a sense that each customer j buys all items in e_j if the sum of the prices does not exceed her budget (i.e. $b_j \geqslant \sum_{i \in e_j} p(i)$, where $p(i)$ indicates the price of item i), in which the seller gets $\sum_{i \in e_j} p(i)$ from the customer. Otherwise, the customer does not buy anything and the seller gets no profit from this customer. The goal of the seller is to set a nonnegative price to each item to maximize her profit from m customers.

This problem was proposed by Guruswami et al. [15], and has received much attention. Let k be the maximum cardinality of any e_i. Approximability of this problem achieved by polynomial time algorithms for large k and n is relatively well-understood now. There is a polynomial time algorithm that guarantees $O(\min(k, (n \log n)^{1/2}))$ fraction of the optimal solution, while we cannot hope for an approximation ratio better than $\Omega(\min(k^{1-\epsilon}, n^{1/2-\epsilon}))$ for any $\epsilon > 0$ under the Exponential Time Hypothesis [6].

The special case $k = 2$ has also been studied in many works separately. The instance can be nicely represented by a graph, with vertices as items and edges as customers, so this problem is called the *Graph (Vertex) Pricing* problem. The fact that this case can be represented as a graph not only gives a theoretical simplification, but also makes the problem flexible to model other settings. For example, Lee et al. [25, 26] independently suggested the same problem from the networking community, motivated by the study of pricing traffic between different levels of internet service providers under the presence of peering.

The best known approximation algorithm for a general instance of Graph Pricing, which guarantees $\frac{1}{4}$ of the optimal solution, is given by Balcan and Blum [4] and Lee et al. [25] The algorithm is simple enough to state here. First, assign 0 to each vertex with probability half independently. For each remaining vertex v, assign the price which maximizes the profit between v and its neighbors already assigned 0. This simple algorithm has been neither improved nor proved to be optimal. Graph Pricing is APX-hard [15], but the only strong hardness of approximation result rules out an approximation algorithm with a guarantee better than $\frac{1}{2}$ [19] under the Unique Games Conjecture (UGC) (via reduction from Maximum Acyclic Subgraph).

The $\frac{1}{4}$-approximation algorithm is surprisingly simple and does not even rely on the power of a linear programming (LP) or semidefinite programming (SDP) relaxation. The efforts to exploit the power of LP relaxations to find a better approximation algorithm have produced positive results for special classes of graphs. Krauthgamer et al. [23] studied the case where all budgets are the same (but the graph might have a self-loop), and proposed a $\frac{6+\sqrt{2}}{5+\sqrt{2}} \approx 1.15$-approximation algorithm based on a LP relaxation. In general case, the standard LP is shown to have an integrality gap close to $\frac{1}{4}$ [19]. Therefore, it is natural to consider hierarchies of LP relaxations such as the Sherali-Adams hierarchy [33] (see [11] for a general survey and [14, 35] for recent algorithmic results using the Sherali-Adams hierarchy). Especially, Chalermsook et al. [5] recently showed that there is a FPTAS when the graph has bounded treewidth, based on the Sherali-Adams hierarchy. However, the power of the Sherali-Adams hierarchy and SDP, as well as the inherent hardness of the problem, was not well-understood in general case.

1.1 Our Results

In this work, we show that any polynomial time algorithm that guarantees a ratio better than $\frac{1}{4}$ must be powerful enough to refute the Unique Games Conjecture.

Theorem 1.1. *Under the Unique Games Conjecture, for any $\epsilon > 0$, it is NP-hard to approximate Graph Pricing within a factor of $\frac{1}{4} + \epsilon$.*

By the results of Khot and Vishnoi [22] and Raghavendra and Steurer [31] that convert a hardness under the UGC to a SDP gap instance, our result unconditionally shows that even a SDP-based algorithm will not improve the performance of a simple algorithm. For the Sherali-Adams hierarchy, we prove that even polynomial rounds of the Sherali-Adams hierarchy has an integrality gap close to $\frac{1}{4}$.

Theorem 1.2. *Fix $\epsilon > 0$. There exists $\delta > 0$ such that the integrality gap of n^δ-rounds of the Sherali-Adams hierarchy for Graph Pricing is at most $\frac{1}{4} + \epsilon$.*

Our result is based on an interesting generalization of Max-Dicut, which we call Generalized Max-Dicut. It is parameterized by a positive integer $T \geqslant 1$. An instance consists of a directed graph $D = (V, A)$ and a label on each edge $l_A : E \to \{1, \ldots, T\}$, where the goal is to assign to each vertex v a label $l_V(v)$ from $\{0, \ldots, T\}$ to maximize the number of satisfied edges — each edge (u, v) is satisfied if $l_V(u) = 0$ and $l_V(v) = l_A(u, v)$.

This problem shares many properties with Graph Pricing, including a simple combinatorial $\frac{1}{4}$-approximation algorithm. There is an approximation-preserving reduction from Generalized Max-Dicut(T) on directed acyclic graphs (DAGs) to Graph Pricing for any T. We prove the following theorems that it is hard to improve upon this simple algorithm for large T even on DAGs, which immediately imply Theorem 1.1 and 1.2.

Theorem 1.3. *Under the Unique Games Conjecture, it is NP-hard to approximate Generalized Max-Dicut(T) on directed acyclic graphs within a factor of $\frac{1}{4} + O(\frac{1}{T^{1/4}})$.*

Theorem 1.4. *Fix T and $\epsilon > 0$. There exists $\delta > 0$ such that the integrality gap of n^δ-rounds of the Sherali-Adams for Generalized Max-Dicut(T) is at most $\frac{T+1}{4T}(1 + \epsilon)$. Furthermore, the same result holds even when the graph is acyclic.*

It is also interesting to compare the above results to other arity two Constraint Satisfaction Problems (CSPs), since whether the domain is Boolean (e.g. Max-Cut, Max-2SAT [13]) or not (e.g. 2-CSP with bounded domain [18], Unique Games [8]), SDP-based algorithms give a strictly better guarantee than LP-based or combinatorial algorithms. As discussed above, our result unconditionally says that a SDP-based algorithm cannot outperform a simple combinatorial algorithm for this arity two CSP (as T increases).[1]

1.2 Related Work and Our Techniques

Formulation of Generalized Max-Dicut Our conceptual contribution is the introduction of Generalized Max-Dicut as a CSP that captures the complexity of Graph Pricing. It is inspired by the work of Khandekar et al. [19], and our reduction is the almost same as their reduction from Max-Acyclic Subgraph (MAS) to Graph Pricing.

In the natural formulation of Graph Pricing as a CSP, each vertex is assigned an (half-)integer price from 0 to B for the maximum budget B, and each customer becomes multiple constraints on two variables since the payoff linearly depends on the prices. It is shown in [19] that a half-integral optimal solution always exists for integral budgets, so this is a (almost) valid relaxation. However, as each customer becomes multiple constraints with different payoffs, it seems hard to apply current techniques developed for well-studied CSPs to this formulation.

[1] Formally, (approximation ratio of the SDP-based algorithm) / (approximation ratio of the best known combinatorial algorithm) = $1 + O(\frac{1}{T^{1/4}})$ for Generalized Max-Dicut. For Unique Games with T labels, the SDP-based algorithm of Charikar et al. [8], which satisfies roughly $T^{-\epsilon/(2-\epsilon)}$ fraction of constraints in an $(1 - \epsilon)$-satisfiable instance, performs better than the random assignment by any constant factor as T increases.

Khandekar et al.'s main idea was to use two well-known CSPs — MAS for the hardness of approximation and Max-Dicut on directed acyclic graphs for the integrality gap of the standard LP. The former is harder to approximate, and the latter has the lower optimum.[2] Generalized Max-Dicut seems to combine ingredients of both problems needed for Graph Pricing. It certainly inherits properties of Max-Dicut including low integral optima, but is much harder to approximate than Max-Dicut by Theorem 1.3.

Uniques Games-Hardness Proving hardness of Generalized Max-Dicut on general graphs is relatively straightforward — proposing a dictatorship test with high completeness and low soundness, and plugging it into the *recipe* of Khot et al. [21] to deduce the hardness result. The dictatorship test is an instance of Generalized Max-Dicut with the set of vertices $\{0, ..., T\}^R$ (called *hypercube*) for some $R \in \mathbb{N}$. The main question in constructing a dictatorship test is how to sample $(x, y) \in \{0, \dots, T\}^2$, which induces a distribution on $\{0, \dots, T\}^2$. In Generalized Max-Dicut, 0 is the only special label such that every directed edge is satisfied only if its tail is assigned 0. The simple combinatorial algorithm samples 0 heavily — the marginal distribution satisfies $\Pr[x = 0] \geqslant 0.5$, while the solution to the Sherali-Adams hierarchy constructed in Theorem 1.4 treats 0 as other labels, having $\Pr[x = 0] = \frac{1}{T+1}$. The latter distribution had a disadvantage that x and y are perfectly correlated — the value of x determines the value of y.

To show the hardness based on the UGC (roughly equivalent to constructing a solution that *fools* SDP), we found that $\Pr[x = 0] = \frac{1}{T^{1/4}}$ is enough. In this case, we can ensure that the probability that dictators pass the test is large, while x and y behave almost independently. Based on the low correlation, we use the result of Mossel [28] to show low soundness.

The resulting dictatorship test is not a DAG. To fix this problem, the final dictatorship test has the vertex set $V \times [T]^R$ for some DAG $D = (V, A)$. For each edge $(u, v) \in A$, the above dictatorship test is performed so that each edge of the dictatorship test goes from the hypercube associated with u to the one with v. This idea of keeping the dictatorship test acyclic is used in Svensson [34], where he takes (the undirected version of) D to be a complete graph. We take a nontrivial DAG found by Alon et al. [1] where any directed cut has at most $\frac{1}{4} + o(1)$ fraction of edges. In the soundness case, if every hypercube is *pseudorandom*, the soundness analysis of an individual dictatorship test associated with each edge gives a rounding algorithm that finds a large directed cut in D, which contradicts the choice of D.

This style of argument, composing the dictatorship test with a certain instance and solving this instance by the soundness analysis, resembles that of Raghavendra [30] for CSPs, Guruswami et al. [16] for ordering CSPs, Kumar et al. [24] for strict CSPs, and Guruswami and Saket [17] for k-uniform k-partite Hypergraph Vertex Cover. While they require the instance to have a good fractional solution (LP or SDP) but the low integral optimum, we only need the low integral optimum (of even a simpler problem) and our individual dictatorship test ensures completeness and part of soundness. We hope that this two-level technique — constructing a simple dictatorship test for each edge and composing it with a certain instance with purely combinatorial properties — makes it easier to bypass the barrier of finding a gap instance and prove hardness for many other problems, especially those with structured instances.

Sherali-Adams Gap On the integrality gap of Generalized Max-Dicut on a DAG, our work generalizes the work of Charikar et al. [9], which showed a similar result for Max-Cut, in several directions. The first obstacle is to find a DAG with a low integral optimum which is amenable to construct a good solution to the Sherali-Adams hierarchy. Previous works which obtained lower bounds for the Sherali-Adams hierarchy [2, 12, 9] used $G(n, p)$, but $G(n, p)$ with an consistent orientation will not result in a low integral optimum.

[2] Under the Unique Games Conjecture, the best inapproximability ratio is 0.5 for MAS [16] and 0.874 for Max-Dicut [3]. For the lower bound on integral optima, the maximum acyclic subgraph always has at least half of edges, while there is a directed acyclic graph where every directed cut cannot have more than $\frac{1}{4} + \epsilon$ fraction of edges for any $\epsilon > 0$.

Instead, we show that sparsifying the aforementioned graph constructed in Alon et al. [1], which is already a DAG with a low integral optimum, gives other desired properties as well.

Given a set S of k vertices, we define a local distribution on the events $\{l_V(v) = i\}_{v \in S, i \in T}$. One caveat of the above approach is that local distributions obtained might be inconsistent, in a sense that S and S' might induce different marginal distributions on $S \cap S'$. Charikar et al.'s main idea is to embed them into l_2 and use hyperplane rounding to produce consistent ones. The most technical part of our work is to extend the hyperplane rounding to work for non-Boolean domains. It is a complicated task in general, but we use the fact that the embedding is explicitly constructed for two adjacent vertices and it exhibits some symmetry, so that we can analyze the performance of our rounding. For $T = 1$, our result matches that of [9].

1.3 Organization

Section 2 introduces problems and notations formally. Section 4 and Section 5 present Unique Games-hardness and Sherali-Adams integrality gaps of Generalized Max-Dicut respectively, which can be combined with the reduction in Section 3 to give the same results for Graph Pricing.

2 Preliminaries

For any positive integer n, let $[n] := \{0, 1, 2, \ldots, n\}$ and $[n]^+ := \{1, 2, \ldots, n\}$. Given a sequence of numbers $a_1, ..., a_n$, let $\max 2_j[a_j]$ be the second largest number among a_j's.

Graph Pricing An instance of Graph Pricing consists of an undirected (possibly contain parallel edges) graph $G = (V, E)$ with budgets $b : E \to \mathbb{R}^+$ and weights $w : E \to \mathbb{R}^+$. Our goal is to find a pricing $p : V \to \mathbb{R}^+ \cup \{0\}$ to maximize

$$\mathsf{Val}(p) := \sum_{e=(u,v) \in E} w(e)(p(u) + p(v))\mathbb{I}[p(u) + p(v) \leqslant b(e)]$$

where $\mathbb{I}[\cdot]$ is the indicator function. Let $\mathsf{Opt}(G, b, w) := \max_p \mathsf{Val}(p)$.

Remark 2.1. *This definition of Graph Pricing above coincides with General Graph Pricing defined in Khandekar et al. [19]. They presented an additional reduction from General Graph Pricing to Graph Pricing with no parallel edge and $w(e) = 1$. Throughout this paper, we use the definition above and allow weights and parallel edges for simplicity. In practice, weights can be naturally interpreted as the number of customers interested in the same pair.*

Remark 2.2. *Another well-known pricing problem assumes that each customer will buy the cheapest item of her interest if she can afford it (i.e., $\mathsf{Val}(p) := \sum_{e=(u,v)} w(e) \min(p(u), p(v))\mathbb{I}[\min(p(u), p(v)) \leqslant b(e)]$), which is called* unit-demand pricing. *Its approximability is similar to that of our single-minded pricing, including algorithms / hardness results for k-Hypergraph Pricing for large k [6], and a simple $\frac{1}{4}$-approximation algorithm for Graph Pricing ($k = 2$). Indeed, Generalized Max-Dicut is also reducible to Unit-demand Graph Pricing and Theorem 1.1 and 1.2 hold for it as well. We focus on Single-minded Graph Pricing here.*

Generalized Max-Dicut Fix a positive integer T. An instance of Generalized Max-Dicut(T) consists of a digraph $D = (V, A)$ with a label $l_A : A \to [T]^+$ and a weight $w : A \to \mathbb{R}^+$ on each edge. Assume that the sum of weights is normalized to 1. (u, v) denotes the edge of D from u to v. We allow parallel edges

from u to v if they have different labels (if parallel edges have the same label, simply merge them). Our goal is to find a labeling $l_V : V \to [T]$ (note vertices can be assigned 0, while edges are not) to maximize the weight of *satisfied* edges — (u, v) is satisfied when $l_V(u) = 0$ and $l_V(v) = l_A(u, v)$. Note than when $T = 1$, the problem becomes Max-Dicut. Given an instance $D = (V, A)$, l_A, and w, let $\mathsf{Opt}(G, l_A, w)$ be the maximum weight of edges satisfied by any labeling of vertices. Given an assignment $l_V : V \to [T]$ to the vertices, let $\mathsf{Val}(l_V)$ be the weight of edges satisfied by l_V. Note that unlike Graph Pricing, the value of any assignment is normalized between 0 and 1. The *normalized outdegree*, denoted by ndeg, is defined to be $[\sum_u (\max_{(u,v)\in A} w(u, v))]^{-1}$. In unweighted instances (i.e. $w(e) = \frac{1}{|A|}$ for all e), $\mathsf{ndeg} \geqslant \frac{|A|}{|V|}$.

Sherali-Adams Hierarchy In its most intuitive and redundant form, a feasible solution to the r-rounds of the Sherali-Adams hierarchy for a CSP with the domain $[q]$ consists of $\sum_{i=1}^{r} \binom{n}{r}(q+1)^r$ variables $\{x_S(\alpha)\}$ for each subset of variables S with cardinality at most r, and $\alpha \in [q]^S$. Each $x_S(\alpha)$ can be interpreted as the probability that the variables in S are assigned α. Therefore, it is required to satisfy the following natural conditions: (1) $x_S(\alpha) \geqslant 0$ for all S, α. (2) $\sum_{\alpha\in[q]^S} x_S(\alpha) = 1$ for all S. (3) $\sum_{\alpha\in[q]^{S'\setminus S}} x_{S'}(\alpha \circ \beta) = x_S(\beta)$ for all $S \subseteq S', \beta \in [q]^S$, where $\alpha \circ \beta \in [q]^{S'\cup S}$ denote the joint assignment to the variables in S'.

The r-rounds of the Sherali-Adams hierarchy for Graph Pricing and Generalized Max-Dicut(T) can be obtained by choosing an appropriate domain and an objective function, while using the constraints given above. For Graph Pricing, if we choose the domain to be $[B]$ where B is the maximum budget, the objective function is the following.

$$\sum_{e=(u,v)} w(e) \sum_{(i,j)\in[B]^2, i+j\leqslant b(u,v)} (i + j) \cdot x_{(u,v)}(i, j)$$

Since $p(v)$ can be real, it is not clear whether this is a relaxation, even when the budgets are integers. [19] shows that there is a half-integral optimal solution. The maximum budget B can be exponentially big in the size of an instance, and a standard trick is to consider only the powers of $(1 + \epsilon)$ as valid prices. It loses at most ϵ fraction of the optimum. Our gap instance and proposed solution to the hierarchy have the marginal on each vertex supported by a constant number of prices, so they are applicable to any choice of the domain.

For Generalized Max-Dicut(T), the domain is $[T]$, and the objective function is

$$\sum_{(u,v)\in A} w(u, v) x_{(u,v)}(0, l_A(u, v)).$$

Given an instance and a relaxation, we define the *integrality gap* to be the integral optimum divided by the value of the best solution to the relaxation. Since both our problems are maximization problems, it is at most 1 and a small number indicates a large gap.

3 Reduction from Generalized Max-Dicut to Graph Pricing

Theorem 3.1. *For any $T > 0$, there is a polynomial time reduction from an instance $(D = (V, A), l_A, w_{\mathsf{GMD}})$ of Generalized Max-Dicut(T), where D is acyclic and $\mathsf{ndeg} \geqslant \frac{1}{\epsilon}$, to an instance (G, b, w_{GP}) of Graph Pricing such that $\mathsf{Opt}(D, l_A, w_{\mathsf{GMD}}) \leqslant \mathsf{Opt}(G, b, w_{\mathsf{GP}}) \leqslant \mathsf{Opt}(D, l_A, w_{\mathsf{GMD}}) + 3\epsilon$.*

Proof. Fix an instance $(D = (V, A), l_A, w_{\mathsf{GMD}})$ of Generalized Max-Dicut(T) with $n = |V|$ and $m = |A|$. Let G be the underlying undirected graph of D. Our reduction from Generalized Max-Dicut on directed acyclic graphs to Graph Pricing is almost the same as the one in Khandekar et al. [19] with some simplification. Let M be a large number which will be fixed later.

The resulting instance of Graph Pricing is based on the same graph G. Since D is acyclic, there is an injective function $s : V \to [n]^+$ such that for each edge $(u, v) \in A$, $s(u) > s(v)$. For each edge $(u, v) \in A$, $b(u, v) = M^{Ts(v)+l_A(u,v)-1}$ and $w_{\mathsf{GP}}(u, v) = \frac{w_{\mathsf{GMD}}(u,v)}{b(u,v)}$.

To avoid confusion, let $\mathsf{Opt}_{\mathsf{GMD}}$, $\mathsf{Val}_{\mathsf{GMD}}$ denote Opt, Val for Generalized Max-Dicut instances, and $\mathsf{Opt}_{\mathsf{GP}}$ and $\mathsf{Val}_{\mathsf{GP}}$ for Graph Pricing instances. Fix a labeling $l_V : V \to [T]$. The corresponding *canonical solution* $p : V \to \mathbb{R}^+ \cup \{0\}$ defined by

$$p(v) = \begin{cases} M^{Ts(v)+l_V(v)-1} & \text{if } l_V(v) \neq 0 \\ 0 & \text{otherwise} \end{cases}$$

gives $\mathsf{Val}_{\mathsf{GP}}(p) \geqslant \mathsf{Val}_{\mathsf{GMD}}(l_V)$ — for each $(u, v) \in A$ satisfied by l_V, p gets $p(v)w_{\mathsf{GP}}(u, v) = w_{\mathsf{GMD}}(u, v)$. Therefore, $\mathsf{Opt}_{\mathsf{GP}}(G, b, w_{\mathsf{GP}}) \geqslant \mathsf{Opt}_{\mathsf{GMD}}(D, l_A, w_{\mathsf{GMD}})$. The following lemma shows that the converse is almost true. The proof is given in Appendix A.

Lemma 3.2 ([19])**.** *For any p, $\mathsf{Val}_{\mathsf{GP}}(p) \leqslant \mathsf{Opt}_{\mathsf{GMD}}(D, l_A, w_{\mathsf{GMD}}) + \frac{1}{M} + 2\epsilon$.*

Taking $M \geqslant \frac{1}{\epsilon}$ proves the theorem. □

4 Approximability of Generalized Max-Dicut

Recall that Generalized Max-Dicut(1) is exactly the well-known Max-Dicut problem, which admits a 0.874-approximation algorithm [27] as any Max-2CSP over the Boolean domain. As T increases, however, the best approximation ratio for Max-2CSP over the domain of size $T + 1$ can be at most $O(\frac{\log T}{\sqrt{T}})$ [7], so viewing it as a general Max-2CSP does not yield a constant-factor approximation algorithm.

There is a simple $\frac{1}{4}$-approximation algorithm, similar to the one for Graph Pricing — assign 0 to each vertex with probability half independently and assign nonzero values to the remaining vertices greedily. The proof is based on the fact that we can easily find the optimal solution once the set of vertices assigned 0 is given. For small T, we can do a little better based on a standard LP relaxation. The proof is given in Appendix C.

Theorem 4.1. *There is a polynomial time approximation algorithm for Generalized Max-Dicut(T) that guarantees $\frac{1}{4} + \Omega(\frac{1}{T})$ of the optimal solution.*

However, we prove that for large T, it is Unique Games-hard to improve the approximation ratio from $\frac{1}{4}$ to a better constant.

Theorem 4.2 (Restatement of Theorem 1.3)**.** *Under the Unique Games Conjecture, it is NP-hard to approximate Generalized Max-Dicut(T) on directed acyclic graphs within a factor of $\frac{1}{4} + O(\frac{1}{T^{1/4}})$.*

Together with the reduction shown in Theorem 3.1, it immediately implies Theorem 1.1 for Graph Pricing. Besides working on DAGs, the reduction also requires that ndeg be large, but it can be easily ensured by taking an Unique Games instance with large degree. See Appendix D.2 to see the full details.

The theorem is proved by proposing a *dictatorship test* with high completeness and low soundness, combined with the standard technique to convert a dictatorship test to a hardness result based on the Unique Games Conjecture [21]. Constructing the dictatorship test has two components — a simple dictatorship test based on correlation and Gaussian geometry, and composing it with a designated DAG. We present the dictatorship test here and defer the full reduction from Unique Games to Appendix D.2.

4.1 Dictatorship Test

We follow the notations in Mossel [28]. Consider the hypercube $[T]^R$ where $[T] = \{0, 1, \ldots, T\}$. Let $\Omega_1 = \Omega_2 = [T]$. For $t \in [T]^+$, $\mathbf{P}^t$ is a probability measure on $\Omega_1 \times \Omega_2$. Let $\mathbf{P}$ be the marginal on Ω_i in $\mathbf{P}^t$ (which does not depend on t and i). We want to ensure that $\mathbf{P}(0) = \delta$, $\mathbf{P}(j) = \frac{1-\delta}{T}$ for $j \in [T]^+$ where $\delta = \frac{1}{T^{1/4}}$. Let $\mathbf{P}'$ be the distribution on Ω_1 such that $\mathbf{P}'(0) = (\frac{1}{1-\frac{1-\delta}{T}})(\delta - \frac{1-\delta}{T})$, $\mathbf{P}'(j) = (\frac{1}{1-\frac{1-\delta}{T}})(\frac{1-\delta}{T})$ (subtract $\frac{1-\delta}{T}$ from $\mathbf{P}(0)$ and renormalize). $\mathbf{P}^t$ is defined by the following procedure to sample (x, y). Sample y according to $\mathbf{P}$. If $y = t$, set $x = 0$. Otherwise, sample x from $\mathbf{P}'$ independently. It is easy to see that the marginal of both x and y is $\mathbf{P}$. We show that (x, y) are almost independent as T increases. We define the correlation between two correlated spaces and prove the following lemma in Appendix D.

Definition 4.3. *Given a distribution $\mathbf{Q}$ on $\Omega_1 \times \Omega_2$, we define the correlation $\rho(\Omega_1, \Omega_2; \mathbf{Q})$ by letting*

$$\rho(\Omega_1, \Omega_2; \mathbf{Q}) = \sup \{\mathsf{Cov}[f, g] : f : \Omega_1 \to \mathbb{R}, g : \Omega_2 \to \mathbb{R}, \mathsf{Var}[f] = \mathsf{Var}[g] = 1\}.$$

Lemma 4.4. *For any t, $\rho(\Omega_1, \Omega_2; \mathbf{P}^t) \leqslant \sqrt{\frac{2}{T\delta}}$.*

Another component of the dictatorship test is the directed acyclic graph $D = (V, A)$ of Alon et al. [1], where every directed cut has size at most $(\frac{1}{4} + o(1))|A|$. Fix a graph $D = (V, A)$ such that every dicut cuts at most $(\frac{1}{4} + \frac{1}{T^{1/4}})|A|$ edges. Note that the size of this graph depends only on T. We now describe the dictatorship test. The prover is expected to provide $F_v : [T]^R \to [T]$ for each $v \in V$.

1. Choose $(u, v) \in A$ and $t \in [T]^+$ uniformly at random.
2. For each $i \in [R]^+$, pick (x_i, y_i) according to $\mathbf{P}^t$.
3. Accept if $F_u(x) = 0$ and $F_v(y) = t$.

This dictatorship test can be naturally interpreted as an instance of Generalized Max-Dicut(T) with the vertex set $V \times [T]^R$. The weight of edge $((u, x), (v, y))$ with label t is equal to the probability that it is sampled, and a labeling $l : V \times [T]^R \mapsto [T]$ passes with probability $\mathsf{Val}(l)$ (by $F_v(x) = l(v, x)$).

4.2 Completeness and Soundness

The ith dictator function is $D_i : [T]^R \to [T]$ given by $D_i(x_1, \ldots, x_R) = x_i$. The purpose of the above dictatorship test is to allow dictatorship functions to be accepted with high probability while penalizing functions *far from* any dictator. The following lemma for completeness is immediate from the test — for any fixed t and i, $\Pr[x_i = 0, y_i = t] = \Pr[y_i = t] = \frac{1-\delta}{T}$.

Lemma 4.5 (Completeness)**.** *Suppose that for some i, $F_v = D_i$ for all $v \in V$. The above test accepts with probability $\frac{1-\delta}{T}$.*

For $v \in V$ and $t \in [T]$, let $F_{v,t} : [T]^R \to \{0, 1\}$ be defined such that $F_{v,t}(x) = 1$ iff $F_v(x) = t$, and $\mu_{v,t} := \Pr[F_v(x) = t] = \mathbb{E}[F_{v,t}(x)]$ where $x \sim \mathbf{P}$. For each $F_{v,t}$ and $i \in [R]^+$, we define the *influence* of the ith coordinate to measure the extent a function depends on the ith coordinate.

$$\mathsf{Inf}_i[F_{v,t}] := \mathbb{E}[\mathsf{Var}[F_{v,t}(X_1, \ldots, X_R) | X_j, 1 \leqslant j \leqslant R, j \neq i]].$$

We use a similarly defined *low-degree influence* $\mathsf{Inf}_i^{\leqslant d}$ for our soundness (see [28] for the definition).

Lemma 4.6 (Soundness)**.** *For large enough T, there exist τ and d (depending on T) such that if $\mathsf{Inf}_i^{\leqslant d}(F_{v,t}) \leqslant \tau$ for all $i \in [R]^+$, $t \in [T]$, and $v \in V$, the probability of accepting is at most $\frac{1}{4T} + \frac{4}{T^{5/4}}$.*

Proof. We use the following theorem of Mossel [28].

Theorem 4.7 (Theorem 6.3 of [28])**.** *Let $(\Omega_1 \times \Omega_2, \mathbf{P})$ be correlated spaces such that the minimum nonzero probability of any atom in $\Omega_1 \times \Omega_2$ is at least α and such that $\rho(\Omega_1, \Omega_2; \mathbf{P}) \leqslant \rho$. Then for every $\epsilon > 0$ there exist τ, d depending on ϵ and α such that if $f : \Omega_1^R \to [0,1], g : \Omega_2^R \to [0,1]$ satisfy $\max(\mathsf{Inf}_i^{\leqslant d}[f], \mathsf{Inf}_i^{\leqslant d}[g]) \leqslant \tau$ for all i, then*

$$\mathop{\mathbb{E}}_{(x,y)\in\mathbf{P}^{\otimes R}}[f(x)g(y)] \leqslant \Gamma_\rho(\mathop{\mathbb{E}}_x[f], \mathop{\mathbb{E}}_y[g]) + \epsilon.$$

The probability of accepting is at most

$$\mathop{\mathbb{E}}_{(u,v)\in A}[\mathop{\mathbb{E}}_{t\in[T]^+}[\mathop{\mathbb{E}}_{(x,y)\sim(\mathbf{P}^t)^{\otimes R}}[F_{u,0}(x)F_{v,t}(y)]]] \leqslant \mathop{\mathbb{E}}_{(u,v)\in A}[\mathop{\mathbb{E}}_{t\in[T]^+}[\Gamma_\rho(\mu_{u,0}, \mu_{v,t}) + \frac{1}{T^{5/4}}]]$$

where the inequality follows from Theorem 4.7 (set $\epsilon \leftarrow \frac{1}{T^{5/4}}$ and $\alpha = \Theta(\frac{1}{T^2})$). The following lemma, whose proof is given in Appendix E, shows that it is at most

$$\mathop{\mathbb{E}}_{(u,v)\in A}[\Gamma_\rho(\mu_{u,0}, \frac{1-\mu_{v,0}}{T})] + \frac{1}{T^{5/4}}.$$

Lemma 4.8. *Fix $\rho, a \in (0,1)$. The function $f(b) := \Gamma_\rho(a,b)$ is concave.*

The following lemma, whose proof is again given in Appendix E, shows that it is at most

$$\mathop{\mathbb{E}}_{(u,v)\in A}[\mu_{u,0}\frac{(1-\mu_{v,0})}{T} + \frac{2}{T^{5/4}}] + \frac{1}{T^{5/4}} = \frac{1}{T}\mathop{\mathbb{E}}_{(u,v)\in A}[\mu_{u,0}(1-\mu_{v,0})] + \frac{3}{T^{5/4}}.$$

Lemma 4.9. *For large enough T and $\delta = \frac{1}{T^{1/4}}$, the following holds. For any $a \in [0,1], b \in [0, \frac{1}{T}]$ and $\rho \in (0, \sqrt{\frac{2}{T\delta}})$, $\Gamma_\rho(a,b) \leqslant ab + \frac{2}{T^{5/4}}$.*

Given $\{\mu_{v,0}\}_{v\in V}$, imagine the rounding algorithm which puts $v \in S$ with probability $\mu_{v,0}$ independently. The expected fraction of edges from S to $V\setminus S$ is $\mathbb{E}_{(u,v)\in A}[\mu_{u,0}(1-\mu_{v,0})]$, which is at most the fractional size of maximum dicut of D. Since we took D to satisfy that $\mathbb{E}_{(u,v)\in A}[\mu_{u,0}(1-\mu_{v,0})] \leqslant \frac{1}{4} + \frac{1}{T^{1/4}}$, the probability of accepting is at most $\frac{1}{4T} + \frac{4}{T^{5/4}}$ as desired. Note that the probabilities of accepting in completeness and soundness differ by a factor of $\frac{\frac{1}{4T}+\frac{4}{T^{5/4}}}{\frac{1-\delta}{T}} = \frac{\frac{1}{4}+\frac{4}{T^{1/4}}}{1-\frac{1}{T^{1/4}}} = \frac{1}{4} + O(\frac{1}{T^{1/4}})$. ∎

5 Integrality Gaps for Generalized Max-Dicut

Fix a positive integer T and $\epsilon \in (0, \frac{1}{100})$. We present an instance of Generalized Max-Dicut(T) $(D = (V,A), l_A)$ (we only deal with unweighted instances in this section and omit w) such that D is acyclic, $|V| \leqslant \epsilon|A|$ (so that $\mathsf{ndeg} \geqslant \frac{1}{\epsilon}$), and a solution to n^δ-rounds of the Sherali-Adams hierarchy such that the integrality gap is at most $\frac{T+1}{4T}(1+\epsilon)$. This result almost matches a simple $\frac{1}{4}$-approximation algorithm.

Through the reduction given in Theorem 3.1, we also prove Theorem 1.2 — a bad integral solution is guaranteed by the reduction, a good solution to the Sherali-Adams hierarchy is obtained by the mapping $l_V(u) = i$ to $p(u) = M^{Ts(u)+i-1}$ (if $i \neq 0$) or 0 (otherwise). The budget in the resulting instance is an integer exponential in the size of instances, and our gap works even for a strong linear programming hierarchy where there is a variable for each vertex v and an integer price i.

The rest of this section is devoted to the proof of Theorem 1.4.

5.1 Obtaining a Good Instance

Our graph D is obtained by randomly sparsifying the graph $D_* = (V, A_*)$ constructed in Alon et al. [1], followed by an appropriate postprocessing. D_* is a directed acyclic graph with n vertices and $m_* = \Theta(n^{\frac{5}{3}})$ edges. Its underlying undirected graph $G_* = (V, E_*)$ is a simple graph with the same number of vertices and edges, with the maximum degree $\Delta_* = \Theta(n^{\frac{2}{3}})$. Actually, $V = [n]^+$ and $(u, v) \in E$ only if $|u - v| \leqslant r$ where $r := \Theta(n^{\frac{2}{3}})$. It has the property that any directed cut has size at most $\frac{m_*}{4} + o(m_*)$ edges.

The first version of $D = (V, A)$ is constructed as the following. $V := V_* = [n]^+$, and for each edge $(u, v) \in A_*$, put $(u, v) \in A$ with probability $p := \frac{\Delta}{\Delta_*}$ for some Δ to be fixed later. Let $G = (V, E)$ be the underlying undirected graph of D. l_A is obtained by assigning each $l(u, v)$ a random number from $[T]^+$.

Like previous integrality gap constructions for Max-Cut and Min-Vertex Cover (e.g. [2, 12, 32, 9]) , D must be postprocessed to be amenable to have a Sherali-Adams solution with a large value. Intuitively, we need to have the underlying undirected graph G *locally sparse* — if we look at a neighborhood of a certain vertex, the graph almost looks like a tree. We use the notion of [9] to measure how locally sparse the graph is.

Definition 5.1. *We say that G' is l-path decomposable if every 2-connected subgraph H of G' contains a path of length l such that every vertex of the path has degree 2 in H.*

The first version of the instance already has $\mathsf{Opt}(D, l_A) \approx \frac{1}{4T}$ with high probability. In order to make the instance locally sparse, we additionally need to remove some of the edges, but the fraction of removed edges is so small that it does not affect $\mathsf{Opt}(D, l_A)$ too much. As a result, we get the following theorem. The proof is given in Appendix B.

Theorem 5.2. *Given T and $\epsilon, \mu > 0$, there exist constants Δ, δ and $l = \Theta(\log n)$ (all constants depending on T and ϵ, μ) such that there is an instance of Generalized Max-Dicut(T) (D, l_A) with the underlying undirected graph G with the following properties.*

- *Acyclicity: D is a DAG.*
- *Low integral optimum: $\mathsf{Opt}(D, l_A) \leqslant \frac{1+\epsilon}{4T}$.*
- *Almost regularity: Maximum degree of G is at most 2Δ, and G has at least $\Omega(\Delta n)$ edges.*
- *Local sparsity: For $k \leqslant n^\delta$, every induced subgraph of G on $(2\Delta)^l k$ vertices is l-path decomposable.*
- *Large noise: For $k \leqslant n^\delta$, $(1 - \mu)^{l/10} \leqslant \frac{\mu}{5k}$.*

The last condition, large noise, is needed to ensure that in a LP solution, even though adjacent vertices are very correlated to give a large value, far away vertices behave almost independently. The meaning of each condition will be elaborated in later sections.

5.2 Constructing (Inconsistent) Local Distributions

Let $D = (V, A)$, l_A, and $G = (V, E)$ be the instance of Generalized Max-Dicut(T) and its underlying undirected graph constructed as above. In this subsection, given a set of $k \leqslant n^\delta$ vertices $S = \{v_1, \ldots, v_k\}$ we give a distribution on events

$$\{l_V(v_1) = x_1, \ldots, l_V(v_k) = x_k\}_{x_1, \ldots, x_k \in [T]}.$$

The local distributions we construct in this subsection are not consistent; for different sets S and S', the marginal distribution on $S \cap S'$ from the distribution on S can be different from the same marginal from the distribution on S' (albeit they are close). This problem is fixed in the next subsection.

Let $d(u,v)$ be the shortest distance between u and v in G and $V' \subseteq V$ be the set of vertices whose shortest distance to S is at most l. Let G' and D' be the subgraph of G and D induced on V', respectively. Since $|V'| \leqslant (2\Delta)^l k$, G' is l-path decomposable by Theorem 5.2. Note that if $d(u,v) < l$, $d(u,v)$ is also the shortest distance between u and v in G'. By the definition, a l-path decomposable graph does not have a cycle of length l, so if $d(u,v) < \frac{l}{2}$, the shortest path between u and v must be unique.

We begin by establishing a fact that when G' is path-decomposable (intuitively looks similar to a tree), there is a distribution on the partitions of V (i.e. multicuts) such that close vertices are unlikely to be separated but far vertices are likely to be separated. If G' is a tree, it is obtained by deleting each edge independently with probability μ. The noise parameter μ will be fixed later depending only on T and ϵ, so is asymptotically greater than $\frac{1}{l} = O(\frac{1}{\log n})$.

Theorem 5.3 ([10]). *Suppose $G' = (V, E)$ is an l-path decomposable graph. Let $L = \lfloor l/9 \rfloor; \mu \in [1/L, 1]$. Then there exists a probabilistic distribution of multicuts of G' (or in other words random partition of G' in pieces) such that the following properties hold. For every two vertices u and v,*

1. *If $d(u,v) \leqslant L$, then the probability that u and v are separated by the multicut (i.e. lie in different parts) equals $1-(1-\mu)^{d(u,v)}$; moreover, if u and v lie in the same part, then the unique shortest path between u and v also lies in that part.*
2. *If $d(u,v) > L$, then the probability that u and v are separated by the multicut is at least $1-(1-\mu)^L$.*
3. *Every piece of the multicut partition is a tree.*

Based on this random partitioning, we define the distribution on the vertices in S (actually in V'). For each piece which is a tree, pick an arbitrary vertex v in the tree, choose $l_V(v)$ uniformly at random, and propagate this label to *weakly satisfy* every edge in the tree — an undirected edge $(u', v') \in E$ (swap u' and v' if necessary to assume $(u', v') \in A$) is weakly satisfied when $l_V(v') - l_V(u') = l_A(u', v')$ over $\mathbb{Z}_{T+1}$. Note that this definition is necessary for the original definition of satisfaction, but not sufficient.

It is clear that the choice of root in each tree does not matter, and the marginal distribution of each $l_V(v)$ is uniform on $[T]$. For vertices u and v with $d(u,v) \leqslant L$, we say that label i for u and i' for v *match* if $l_V(u) = i, l_V(v) = i'$ can be extended to weakly satisfy every edge on the unique shortest path between u and v (there are $T+1$ such pairs). If u and v are close, $l_V(u)$ and $l_V(v)$ will be correlated in a sense that if i and i' match, $l_V(u) = i$ almost implies $l_V(v) = i'$, while it is not the case when u and v are far apart. The following corollary formalizes this intuition. The proof is in Appendix B.2.

Corollary 5.4. *Suppose $G' = (V', E')$ is an l-path decomposable graph. Let $L = \lfloor l/9 \rfloor$; $\mu \in [1/L, 1]$. Then there exists a random mapping $r : V' \to [T]$ such that*

1. *If $d := d(u,v) \leqslant L$ then*

$$\Pr[r(u) = i, r(v) = i'] = \begin{cases} \frac{(1-\mu)^d}{(T+1)} + \frac{1-(1-\mu)^d}{(T+1)^2} & \text{if } i \text{ and } i' \text{ match} \\ \frac{1-(1-\mu)^d}{(T+1)^2} & \text{otherwise} \end{cases}$$

2. *If $d > L$ then $\frac{1-(1-\mu)^L}{(T+1)^2} \leqslant \Pr[r(u) = i, r(v) = i'] \leqslant \frac{1-(1-\mu)^L}{(T+1)^2} + \frac{(1-\mu)^L}{T+1}$ for any $i, i' \in [T]$.*

Definition 5.5. *For any vertices $u \neq v$ and $i, i' \in [T]$, let $\rho(u(i), v(i')) := \Pr[r(u) = i, r(v) = i']$ if $d(u,v) \leqslant L$, or $\frac{1}{(T+1)^2}$ otherwise. $\rho(v(i), v(i)) := \frac{1}{T+1}$ and $\rho(v(i), v(i')) := 0$ for $i \neq i'$. Since the shortest path between u and v is unique when $d(u,v) \leqslant L$, ρ is uniquely defined given G, D, l_A and does not depend on S, V', G', D' which induce a local distribution.*

Definition 5.6. *Fix a set of k vertices $S = \{v_1, \ldots, v_k\}$. For any vertex $u, v \in S$ and $i, j \in [T]$, let $\nu_S(u(i), v(i')) := \Pr[x(u) = i, x(v) = i']$ in the local distribution on S defined by r in Corollary 5.4.*

5.3 Geometric Embedding and Rounding

In this subsection, we still fix a set of k vertices $S = \{v_1, \ldots, v_k\}$ and produce a distribution on the events $\{l_V(v_1) = x_1, \ldots, l_V(v_k) = x_k\}_{x_1,\ldots,x_k \in [T]}$. The difference from the last subsection is that the resulting distributions become consistent — the marginal distribution on $S \cap S'$ does not depend on the choice of its superset (S or S') that is used to obtain a larger local distribution.

5.3.1 Embedding

Consider ρ and ν_S defined in the last subsection. ρ and ν_S both capture the pairwise distribution between the events $\{l_V(v) = x\}_{v \in S, x \in [T]}$, but each of them has its own defects. ν_S depends on the choice of S, so does not yield consistent local distributions. ρ does not depend on S, but for far vertices, Corollary 5.4 does not guarantee any local distribution consistent with it. However, they are close in a sense — they are identical when $d(u, v) \leqslant L$ and differ by at most $\frac{(1-\mu)^L}{T+1}$ otherwise.

The main idea of Charikar et al. [9] is to interpret ρ and ν_S as pairwise distances between events and embed ρ to l_2 with small error. It is based on the fact that ρ and ν_S are close for any S and ν_S is readily embeddable to l_2. Since the embedding into l_2 is uniquely defined by the pairwise distances and ρ does not depend on the choice of S, geometric rounding schemes based on the embedding yield consistent local distributions. Let $v(i)$ be the vector corresponding to the event $l_V(v) = i$. Our goal is to construct $k(T + 1)$ vectors $\{v(i)\}_{v \in S, i \in [T]}$ such that $u(i) \cdot v(i') \approx \rho(u(i), v(i'))$. Following the above intuition, the following lemma says that this embedding is possible with error depending on μ. The proof is given in Appendix B.3.

Lemma 5.7. *There exist $k(T + 1)$ vectors $\{v(i)\}_{v \in S, i \in [T]}$ such that $\|v(i)\|_2^2 = \mu + \frac{1}{T+1}$ and $u(i) \cdot v(i') = \frac{\mu}{2} + \rho(u(i), v(i'))$.*

5.3.2 Rounding and Analyzing adjacent vertices

Given $k(T + 1)$ vectors $\{v(i)\}_{v \in S, i \in [T]}$, our rounding scheme is one of the most natural ways to choose one out of $(T + 1)$ vectors — take a random Gaussian vector g and for each vertex v, set $l_V(v) = i$ such that $v(i) \cdot g$ is the maximum over all i. Since the inner products of these vectors depend only on ρ (which does not depend on the choice of S), it gives a consistent local distribution.

Fix adjacent vertices v and u (without loss of generality assume $(u, v) \in A$). It only remains to show that $\Pr[l_V(u) = 0, l_V(v) = l_A(u, v)] \approx \frac{1}{T+1}$. For any pair of adjacent vertices, we can write $2(T + 1)$ vectors explicitly. They are just two sets of $T + 1$ orthonormal vectors, very closely correlated — there are $T + 1$ pairs $(u(i), v(i'))$, $i' - i = l_A(u, v)$ in $\mathbb{Z}_{T+1}$, such that $u(i) \approx v(i')$. With this symmetric structure and a suitable choice of the noise parameter μ, we can analyze the performance of our rounding. The proof is given in Appendix B.4.

Lemma 5.8. *There exists μ depending on T and ϵ such that, in the above rounding scheme, the probability that $l_V(u) = 0$ and $l_V(v) = l_A(u, v)$ is at most $\frac{1-12\epsilon}{T+1}$.*

This finishes the construction of a solution to the n^δ-rounds of the Sherali-Adams hierarchy with value $\frac{1-12\epsilon}{T+1}$. Since $\mathsf{Opt}(V, l_A) \leqslant \frac{1+\epsilon}{4T}$ by Theorem 5.2, it proves Theorem 1.4 and Theorem 1.2.

Acknowledgements. The author would like to thank Venkat Guruswami and Seung Woo Shin for helpful discussions.

References

[1] N. Alon, B. Bollobás, A. Gyárfás, J. Lehel, and A. Scott. Maximum directed cuts in acyclic digraphs. *Journal of Graph Theory*, 55(1):1–13, 2007. 3, 7, 9, 15

[2] S. Arora, B. Bollobás, L. Lovász, and I. Tourlakis. Proving integrality gaps without knowing the linear program. *Theory of Computing*, 2(1):19–51, 2006. 3, 9, 17

[3] P. Austrin. Towards sharp inapproximability for any 2-CSP. *SIAM Journal on Computing*, 39(6):2430–2463, 2010. 3

[4] M.-F. Balcan and A. Blum. Approximation algorithms and online mechanisms for item pricing. *Theory of Computing*, 3:179–195, 2007. 1

[5] P. Chalermsook, S. Kintali, R. J. Lipton, and D. Nanongkai. Graph pricing problem on bounded treewidth, bounded genus and k-partite graphs. *Chicago Journal of Theoretical Computer Science*, 2013(13), November 2013. 1

[6] P. Chalermsook, B. Laekhanukit, and D. Nanongkai. Independent set, induced matching, and pricing: Connections and tight (subexponential time) approximation hardnesses. In *Proceedings of the 54th annual IEEE symposium on Foundations of Computer Science*, FOCS '13, pages 370–379, 2013. 1, 4

[7] S. O. Chan. Approximation resistance from pairwise independent subgroups. In *Proceedings of the 45th annual ACM Symposium on Theory of Computing*, STOC '13, pages 447–456, 2013. 6

[8] M. Charikar, K. Makarychev, and Y. Makarychev. Near-optimal algorithms for unique games. In *Proceedings of the 38th annual ACM Symposium on Theory of Computing*, STOC '06, pages 205–214, 2006. 2, 25

[9] M. Charikar, K. Makarychev, and Y. Makarychev. Integrality gaps for Sherali-Adams relaxations. In *Proceedings of the 41st annual ACM Symposium on Theory of Computing*, STOC '09, pages 283–292, 2009. 3, 4, 9, 11

[10] M. Charikar, K. Makarychev, and Y. Makarychev. Local global tradeoffs in metric embeddings. *SIAM Journal on Computing*, 39(6):2487–2512, 2010. 10, 18

[11] E. Chlamtac and M. Tulsiani. Convex relaxations and integrality gaps. In *Handbook on Semidefinite, Conic and Polynomial Optimization*, volume 166 of *International Series in Operations Research and Management Science*, pages 139–169. 2012. 1

[12] W. F. de la Vega and C. Kenyon-Mathieu. Linear programming relaxations of maxcut. In *Proceedings of the 18th annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '07, pages 53–61. Society for Industrial and Applied Mathematics, 2007. 3, 9

[13] M. Goemans and D. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM*, 42(6):1115–1145, Nov. 1995. 2

[14] A. Gupta, K. Talwar, and D. Witmer. Sparsest cut on bounded treewidth graphs: Algorithms and hardness results. In *Proceedings of the 45th annual ACM Symposium on Theory of Computing*, STOC '13, pages 281–290, 2013. 1

[15] V. Guruswami, J. D. Hartline, A. R. Karlin, D. Kempe, C. Kenyon, and F. McSherry. On profit-maximizing envy-free pricing. In *Proceedings of the 16th annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '05, pages 1164–1173, 2005. 1

[16] V. Guruswami, J. Håstad, R. Manokaran, P. Raghavendra, and M. Charikar. Beating the random ordering is hard: Every ordering CSP is approximation resistant. *SIAM Journal on Computing*, 40(3):878–914, 2011. 3

[17] V. Guruswami and R. Saket. On the inapproximability of vertex cover on k-partite k-uniform hypergraphs. In *Proceedings of the 37th International Colloquium Conference on Automata, Languages and Programming*, ICALP '10, pages 360–371, 2010. 3

[18] J. Håstad. Every 2-CSP allows nontrivial approximation. *Computational Complexity*, 17(4):549–566, 2008. 2

[19] R. Khandekar, T. Kimbrel, K. Makarychev, and M. Sviridenko. On hardness of pricing items for single-minded bidders. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, volume 5687 of *Lecture Notes in Computer Science*, pages 202–216. 2009. 1, 2, 4, 5, 6, 14

[20] S. Khot. On the power of unique 2-prover 1-round games. In *Proceedings of the 34th annual ACM Symposium on Theory of Computing*, STOC '02, pages 767–775, 2002. 23

[21] S. Khot, G. Kindler, E. Mossel, and R. O'Donnell. Optimal inapproximability results for Max-Cut and other 2-variable CSPs? *SIAM Journal on Computing*, 37(1):319–357, 2007. 3, 6, 24

[22] S. A. Khot and N. K. Vishnoi. The unique games conjecture, integrality gap for cut problems and embeddability of negative type metrics into l_1. In *Proceedings of the 46th annual IEEE symposium on Foundations of Computer Science*, FOCS '05, pages 53–62, 2005. 2

[23] R. Krauthgamer, A. Mehta, and A. Rudra. Pricing commodities. *Theoretical Computer Science*, 412(7):602 – 613, 2011. Selected papers from WAOA 2007: 5th Workshop on Approximation and Online Algorithms. 1

[24] A. Kumar, R. Manokaran, M. Tulsiani, and N. K. Vishnoi. On LP-based approximability for strict CSPs. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '11, pages 1560–1573. SIAM, 2011. 3

[25] E.-w. Lee, D. Buchfuhrer, L. Andrew, A. Tang, and S. Low. Pricing in the presence of peering. In *Proceedings of the 45th annual Allerton Conference on Computing, Communications and Control*, Allerton '07, pages 286–291, 2007. 1

[26] E.-w. Lee, D. Buchfuhrer, L. Andrew, A. Tang, and S. Low. Progress on pricing with peering. In *Proceedings of the 42nd annual Coneference on Information Sciences and Systems*, CISS '08, pages 286–291, 2008. 1

[27] M. Lewin, D. Livnat, and U. Zwick. Improved rounding techniques for the Max 2-Sat and Max Di-Cut problems. In *Proceedings of the 9th international conference on Integer Programming and Combinatorial Optimization*, IPCO '02, pages 67–82, 2002. 6

[28] E. Mossel. Gaussian bounds for noise correlation of functions. *Geometric and Functional Analysis*, 19(6):1713–1756, 2010. 3, 6, 7, 8, 22

[29] R. Otter. The number of trees. *Annals of Mathematics*, pages 583–599, 1948. 17

[30] P. Raghavendra. Optimal algorithms and inapproximability results for every CSP? In *Proceedings of the 40th annual ACM symposium on Theory of computing*, STOC '08, pages 245–254, New York, NY, USA, 2008. ACM. 3

[31] P. Raghavendra and D. Steurer. Integrality gaps for strong SDP relaxations of unique games. In *Proceedings of the 50th annual IEEE symposium on Foundations of Computer Science*, FOCS '09, pages 575–585, 2009. 2

[32] G. Schoenebeck, L. Trevisan, and M. Tulsiani. Tight integrality gaps for Lovasz-Schrijver LP relaxations of vertex cover and max cut. In *Proceedings of the 39th annual ACM Symposium on Theory of Computing*, STOC '07, pages 302–310, 2007. 9

[33] H. Sherali and W. Adams. A hierarchy of relaxations between the continuous and convex hull representations for zero-one programming problems. *SIAM Journal on Discrete Mathematics*, 3(3):411–430, 1990. 1

[34] O. Svensson. Hardness of vertex deletion and project scheduling. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, volume 7408 of *Lecture Notes in Computer Science*, pages 301–312. 2012. 3

[35] Y. Yoshida and Y. Zhou. Approximation schemes via sherali-adams hierarchy for dense constraint satisfaction problems and assignment problems. In *Proceedings of the 5th conference on Innovations in Theoretical Computer Science*, pages 423–438. ACM, 2014. 1

A Proof of Lemma in the Reduction

Lemma A.1 ([19], Restatement of 3.2)**.** *For any p,* $\mathsf{Val}_{\mathsf{GP}}(p) \leqslant \mathsf{Opt}_{\mathsf{GMD}}(D, l_A, w_{\mathsf{GMD}}) + \frac{1}{M} + 2\epsilon$.

Proof. Given p, we define the *principal part* of $\mathsf{Val}_{\mathsf{GP}}(p)$ as

$$\sum_{(u,v)\in A} w_{\mathsf{GP}}(u,v)p(v)\mathbb{I}[p(u)+p(v) \leqslant b(u,v)].$$

Note that for each directed edge, only the price of its head contributes.

We first bound the principal part of $\mathsf{Val}_{\mathsf{GP}}(p)$. For a vertex v, the only edges where $w_{\mathsf{GP}}(u,v)p(v) > \frac{w_{\mathsf{GMD}}(u,v)}{M}$ satisfy $M^{Ts(v)+l_A(u,v)-2} < p(v) \leqslant M^{Ts(v)+l_A(u,v)-1}$. If there is such an edge, let $l'_V(v) = l_A(u,v)$. Otherwise, let $l'_V(v) = 0$. Fix an edge (u,v) where $w_{\mathsf{GP}}(u,v)p(v) > \frac{w_{\mathsf{GMD}}(u,v)}{M}$. $l'_V(v) = l_A(u,v)$ by above. If $l'_V(u) \neq 0$, it means $p(u) > M^{Ts(u)-1} \geqslant M^{Ts(v)+T-1} \geqslant b(u,v)$, so (u,v) contributes 0 to the principal part of $\mathsf{Val}_{\mathsf{GP}}(p)$. Therefore, for each edge (u,v) that contributes more than $\frac{1}{M}$ to the principal part of $\mathsf{Val}_{\mathsf{GP}}(p)$, l'_V satisfies (u,v). Therefore, the principal part of $\mathsf{Val}_{\mathsf{GP}}(p)$ is at most $\mathsf{Opt}_{\mathsf{GMD}}(D, l_A, w_{\mathsf{GMD}}) + \frac{1}{M}$.

For the non-principal part of $\mathsf{Val}_{\mathsf{GP}}(p)$, for each vertex u, and we bound

$$\sum_{(u,v)\in A} w_{\mathsf{GP}}(u,v)p(u)\mathbb{I}[p(u)+p(v) \leqslant b(u,v)] \leqslant \sum_{(u,v)\in A, p(u)\leqslant b(u,v)} w_{\mathsf{GMD}}(u,v)\frac{p(u)}{b(u,v)}.$$

Note that all edges (u, v) have different $b(u, v)$, and any two differ by at least a factor of M. Let $w_u := \max_{(u,v)\in A} w_{\mathsf{GMD}}(u, v)$. Therefore, the right hand side can be bounded by $w_u(1 + \frac{1}{M} + \frac{1}{M^2} + \dots) \leqslant 2w_u$, where

$$\sum_u w_u = \frac{1}{\mathsf{ndeg}} \leqslant \epsilon.$$

This shows that the non-principal part of $\mathsf{Val}_{\mathsf{GP}}(p)$ is at most 2ϵ, proving the lemma. □

B Details of the Integrality Gap

B.1 Obtaining a Good Instance

In this subsection, we prove the following theorem.

Theorem B.1 (Restatement of Theorem 5.2)**.** *Given T and $\epsilon, \mu > 0$, there exist constants Δ, δ and $l = \Theta(\log n)$ (all constants depending on T and ϵ, μ) such that there is an instance of Generalized Max-Dicut(T) (D, l_A) with the underlying undirected graph G with the following properties.*

- *Acyclicity: D is a DAG.*
- *Low Integral Optimum: $\mathsf{Opt}(D, l_A) \leqslant \frac{1+\epsilon}{4T}$.*
- *Almost regularity: Maximum degree of G is at most 2Δ, and G has at least $\Omega(\Delta n)$ edges.*
- *Local Sparsity: For $k \leqslant n^\delta$, every induced subgraph of G on $(2\Delta)^l k$ vertices is l-path decomposable.*
- *Large noise: For $k \leqslant n^\delta$, $(1 - \mu)^{l/10} \leqslant \frac{\mu}{5k}$.*

Proof. As mentioned in Section 5.1, our graph D is obtained by randomly sparsifying the graph $D_* = (V, A_*)$ constructed in Alon et al. [1] after an appropriate postprocessing. D_* is a directed acyclic graph with n vertices and $m_* = \Theta(n^{\frac{5}{3}})$ edges. Its underlying undirected graph $G_* = (V, E_*)$ is a simple graph with the same number of vertices and edges, with the maximum degree $\Delta_* = \Theta(n^{\frac{2}{3}})$. Actually, $V = [n]^+$ and $(u, v) \in E$ only if $|u - v| \leqslant r$ where $r := \Theta(n^{\frac{2}{3}})$. It has the property that any directed cut has size at most $\frac{m_*}{4} + o(m_*)$ edges.

The first version of $D = (V, A)$ is constructed as the following. $V := V_* = [n]^+$, and for each edge $(u, v) \in A_*$, put $(u, v) \in A$ with probability $p := \frac{\Delta}{\Delta_*}$ for some Δ to be fixed later. Let $G = (V, E)$ be the underlying undirected graph of V. l_A is obtained by assigning each $l(u, v)$ a random number uniformly sampled from $[T]^+$.

Integral Solution The following lemma shows that if Δ is big enough, $\mathsf{Opt}(D, l_A)$ is close to $\frac{1}{4T}$.

Lemma B.2. *If G satisfies the above four properties and $\Delta = \Omega(\frac{T \log T}{\epsilon^2})$, then D and l_A obtained by the above process satisfies $\mathsf{Opt}(D, l_A) \leqslant \frac{1+4\epsilon}{4T}$ with high probability.*

Proof. Fix one assignment $l_V : V \to [T]$. For any edge $(u, v) \in A_*$ call it a *candidate* when $l_V(u) = 0, l_V(v) \neq 0$. Note that the number of candidate edges is at most the cardinality of the maximum directed cut of D_*, which is at most $\frac{1+o(1)}{4} m_*$.

For each candidate edge (u, v), the probability that $(u, v) \in A$ with $l_A(u, v) = l_V(v)$ is $\frac{1}{T}$. Therefore, the expected number of satisfied edges is at most $\frac{(1+o(1))\Delta m_*}{4\Delta_* T}$. By Chernoff bound, the probability that it is bigger than $\frac{(1+\epsilon)pm_*}{4T}$ is bounded by $\exp(-\Omega(\frac{\epsilon^2 pm_*}{T})) = \exp(-\Omega(\frac{\epsilon^2 \Delta n}{T}))$. By taking union bound over

$(T+1)^n$ different l_V's, the probability that there exists an assignment with more than $\frac{(1+\epsilon)pm_*}{4T}$ satisfied edges is at most

$$\exp(-\Omega(\frac{\epsilon^2\Delta n}{T})) * \exp(n\log(T+1)) \leqslant n^{-1}$$

for $\Delta := \Omega(\frac{T\log T}{\epsilon^2})$. Similarly, we can conclude that $|A| \geqslant (1-\epsilon)m_*p$ with high probability. Therefore, $\mathsf{Opt}(D, l_A)$ is at most $\frac{(1+\epsilon)}{4T(1-\epsilon)} \leqslant \frac{1+4\epsilon}{4T}$ with high probability. □

The above lemma is the only place where it is desirable to have large $|A| = |E|$. For the rest of this subsection, we are going to delete some edges of D (and G) to satisfy desired properties. Note that in any case, the number of edges deleted is much less than ϵpm_* so that each deletion does not hurt the above lemma.

Maximum Degree Control Since the maximum degree in G_* is Δ_*, expected degree of each vertex $v \in V$ in G is at most $p\Delta_* = \Delta$. Call a vertex $v \in V$ *bad* if it has degree more than 2Δ in G, and call an edge $(u,v) \in E$ bad if either u or v is bad. Fix an edge (u,v). The probability that (u,v) becomes bad given $(u,v) \in E$ is at most $2\exp(-\frac{\Delta}{4})$. The expected number of bad edges is at most $2\exp(-\frac{\Delta}{4})pm_*$, and by Markov's inequality, with probability at least half, the number of bad edges is at most $4\exp(-\frac{\Delta}{3})pm_*$.

Deleting all bad edges guarantees that the maximum degree of G is at most 2Δ, and with probability at least half, we delete only $4\exp(-\frac{\Delta}{3})pm_*$ edges, which is much smaller than ϵpm_* since $\Delta = \Omega(\frac{1}{\epsilon^2})$.

Girth Control The expected number of cycles of length i is bounded by

$$n(2r)^{i-1}p^i = n(2r)^{i-1}(\frac{\Delta}{\Delta^*})^i \leqslant \frac{n(C\Delta)^i}{\Delta_*}$$

for some absolute constant C. When $i = O(\frac{\log n}{\log\Delta})$ the above quantity becomes less than $n^{0.5}$. Assume $l = O(\frac{\log n}{\log\Delta})$ (it will be fixed even smaller than that later). Summing over $i = 4, \dots, l$ ensures that the expected number of cycles of length up to l is at most $O(n^{0.6})$, and it is less than $O(n^{0.7})$ with high probability. Removing one edge for each cycle of length up to l ensures that G has girth at least l.

Local Sparsity Control Let $\eta = \frac{1}{3l}$ for some l fixed later. We want to show that there exists $\gamma > 0$ such that every subgraph G' of G induced on $t \leqslant n^\gamma$ vertices have only $(1+\eta)t$ edges.

For $4 \leqslant t \leqslant 1/\eta$, we count the number of connected subgraphs of G_* with t vertices and $t+1$ edges.

Lemma B.3. *The number of connected subgraphs of G_* with t vertices and $t+1$ edges is bounded by $2nt^2\Delta_*^{t-1}$.*

Proof. The only possible degree sequences for such subgraphs are $(4,2,2,2,\dots)$ or $(3,3,2,2,\dots)$. Assume that it is $(4,2,2,2,\dots)$. Let v be the vertex with degree 4. There is a sequence of $t+2$ vertices $(v,\dots,v,\dots,v)$ representing an Eulerian tour (not necessarily unique). The number of such sequences is bounded by $nt\Delta_*^{t-1}$ (n for guessing v, t for guessing where v occurs in the middle of the sequence, Δ_*^{t-1} for the other vertices).

Assume that the degree sequence is $(3,3,2,2,\dots)$, and u,v be the vertices of degree 3. Take a sequence of $t+2$ vertices representing an Eulerian path from u to v (either $(u,\dots,u,\dots,v,\dots,v)$ or $(u,\dots,v,\dots,u,\dots,v)$). The number of such sequences is bounded by $nt^2\Delta_*^{t-1}$ (n for guessing u, t^2 for guessing positions of u and v in the middle of the sequence, Δ_*^{t-1} for the other vertices including v). □

Therefore, the probability that there exists a subgraph of G with t vertices and $t+1$ edges for $4 \leqslant t \leqslant 1/\eta = 3l$ is

$$\sum_{t=4}^{3l} 2nt^2\Delta_*^{t-1}p^{t+1} = \sum_{t=4}^{3l} \frac{2nt^2\Delta^{t+1}}{\Delta_*^2} \leqslant \frac{n}{\Delta_*^2}(9l)^2\Delta^{3l+1} \leqslant n^{-0.1}$$

for $l = O(\log n/\log\Delta)$, since $\frac{n}{\Delta_*^2} = O(n^{-\frac{1}{3}})$.

For $t > 1/\eta = 3l$, we count the number of subgraphs of G_* with t vertices and $(1+\eta)t$ edges. It is upper bounded by (the number of connected subtrees on t vertices) * (the number of possibilities to choose other $\eta t+1$ edges out of $\binom{t}{2}$ pairs). The number of unlabeled rooted trees on t vertices is $C\alpha^t$ for some constants C and α [29], so the number of connected subtrees on t vertices is bounded by $Cn\alpha^t\Delta_*^{t-1}$. Therefore, the total number of such subgraphs is

$$Cn\alpha^t\Delta_*^{t-1}\binom{\frac{t(t+1)}{2}}{\eta t+1} \leqslant Cn\alpha^t\Delta_*^{t-1}\binom{t^2}{2\eta t} \leqslant Cn\alpha^t\Delta_*^{t-1}(\frac{et}{2\eta})^{2\eta t}.$$

The probability that such a graph exists in G is at most

$$Cn\alpha^t\Delta_*^{t-1}(\frac{et}{2\eta})^{2\eta t}(\frac{\Delta}{\Delta_*})^{(1+\eta)t} \leqslant \frac{n}{\Delta_*}(C_1\Delta^2)^t(C_2\frac{l^2t^2}{\Delta_*})^{t/3l}.$$

Let $A = C_1\Delta^2$ and $B = C_2\frac{l^2t^2}{\Delta_*}$. The above quantity is at most

$$\frac{n}{\Delta_*}A^tB^{t/3l} = (\frac{n}{\Delta_*}A^{3l}B)(AB^{1/3l})^{t-3l}.$$

Assume $t \leqslant n^\gamma$ for some $\gamma \in (0, 0.1)$ and $l = O(\frac{\log n}{\log\Delta})$ be such that $\frac{n}{\Delta_*}A^{3l}B = \frac{C_2l^2t^2n(C_1\Delta^2)^{3l}}{\Delta_*^2} \leqslant n^{-0.1}$, which also implies $AB^{1/3l} \leqslant 1$. Summing over $t = 3l, \ldots, n^\gamma$, the probability that such a graph exists is bounded by $o(1)$.

Putting Them Together In Section 5.1, we mentioned that the resulting graph should be amenable to have a Sherali-Adams solution with a large value, and introduced the notion of path-decomposability to measure it. The following lemma of Arora et al. [2] shows that our construction satisfies that every subgraph of G induced on at most $t \leqslant n^\gamma$ vertices is l-path decomposable.

Lemma B.4 ([2])**.** *Let $l \geqslant 1$ be an integer and $0 < \eta < \frac{1}{3l-1}$, and let H be a 2-connected graph with t vertices and at most $(1+\eta)t$ edges. Then H contains a path of length at least $l+1$ whose internal vertices have degree 2 in H.*

Finally, δ and l are fixed based on the other parameters to satisfy the requirements of the theorem.

Lemma B.5. *There exists $\delta > 0$ and l (depending on T, ϵ, Δ, μ, γ) such that for any $k \leqslant n^\delta$, the following holds.*

1. $(1-\mu)^{\frac{l}{10}} \leqslant \frac{\mu}{5k}$.
2. Every induced subgraph of G on $(2\Delta)^lk$ vertices is l-path decomposable.

Proof. The first condition is implied by $l \geqslant C\delta\log n$ for some constant C depending on μ. The second condition is implied by $(2\Delta)^lk \leqslant n^\gamma \Leftrightarrow l \leqslant C'(\gamma-\delta)\log n$ for another constant C' depending on Δ. When we control girth and local sparsity, l is required to be $O(\frac{\log n}{\log\Delta})$. Therefore, by taking δ a small enough constant depending on T, ϵ, Δ, μ, and γ (all of which depend on T, ϵ), we can ensure that such l exists. □

Therefore, there exist constants Δ, δ and $l = \Theta(\log n)$ (all constants depending on T, ϵ, μ) that satisfy all the requirements given in the theorem. □

B.2 Distribution

Corollary B.6 (Restatement of 5.4)**.** *Suppose $G' = (V', E')$ is an l-path decomposable graph. Let $L = \lfloor l/9 \rfloor$; $\mu \in [1/L, 1]$. Then there exists a random mapping $r : V' \to [T]$ such that*

1. *If $d := d(u, v) \leqslant L$ then*

$$\Pr[r(u) = i, r(v) = i'] = \begin{cases} \frac{(1-\mu)^d}{(T+1)} + \frac{1-(1-\mu)^d}{(T+1)^2} & \text{if } i \text{ and } i' \text{ match} \\ \frac{1-(1-\mu)^d}{(T+1)^2} & \text{otherwise} \end{cases}$$

2. *If $d > L$ then $\frac{1-(1-\mu)^L}{(T+1)^2} \leqslant \Pr[r(u) = i, r(v) = i'] \leqslant \frac{1-(1-\mu)^L}{(T+1)^2} + \frac{(1-\mu)^L}{T+1}$ for any $i, i' \in [T]$.*

Proof. r is defined by the following process: sample a distribution of multicuts as Theorem 5.3. Each piece is a tree, so we can pick an arbitrary vertex w and give a value $l_V(w)$ uniformly from $[T]$ and propagate along the tree to weakly satisfy every edge. Note that the distribution does not depend on the choice of the initial vertex.

Suppose $d(u, v) \leqslant L$, which ensures that if u and v are in the same piece, the only path connecting u and v in the piece is the shortest path in G. If i and i' are match labels,

$$\Pr[r(u) = i, r(v) = i'] = \Pr[u, v \text{ in the same piece}] \cdot \frac{1}{T+1} + \Pr[u, v \text{ separated}] \cdot \frac{1}{(T+1)^2}.$$

If i and i' are nonmatching labels,

$$\Pr[r(u) = i, r(v) = i'] = \Pr[u, v \text{ in the same piece}] \cdot 0 + \Pr[u, v \text{ separated}] \cdot \frac{1}{(T+1)^2}.$$

If $d(u, v) > L$, $\Pr[r(u) = i, r(v) = i']$ is lower bounded by $\frac{\Pr[u \text{ and } v \text{ are separated}]}{(T+1)^2}$, and upper bounded by $\frac{\Pr[u \text{ and } v \text{ are separated}]}{(T+1)^2} + \frac{\Pr[u \text{ and } v \text{ are not separated}]}{T+1}$. The separation guarantee in Theorem 5.3 proves the lemma. □

B.3 Embedding

Lemma B.7 (Restatement of Lemma 5.7)**.** *There exist $k(T+1)$ vectors $\{v(i)\}_{v \in S, i \in [T]}$ such that $\|v(i)\|_2^2 = \mu + \frac{1}{T+1}$ and $u(i) \cdot v(i') = \frac{\mu}{2} + \rho(u(i), v(i'))$.*

Proof. For each $u(i)$, we construct two vectors $u(i)_1$ and $u(i)_2$ and finally merge them by $u(i) := u(i)_1 \oplus u(i)_2$. $u(i)_2$ is the indicator random variable for the event $l_V(u) = i$, where the distribution follows ν_S. Since ν_S is based on an actual distribution on the events, the vectors $\{v(i)_2\}_{v \in V, i \in [T]}$ are embeddable into l_2 with $\|v(i)_2\|_2^2 = \Pr[l_V(v) = i] = \frac{1}{T+1}$ and $u(i)_2 \cdot v(i')_2 = \nu_S(u(i), v(i'))$. The first group of vectors $\{v(i)_1\}_{v \in V, i \in [T]}$ *convert* these inner products from ν_S to ρ with small error.

The following lemma says that a metric space can be isometrically embeddable into l_2 if all pairwise distances are similar.

Lemma B.8 ([10])**.** *Consider a metric space (Y, α) on t points. If for every two distinct points u and v: $|\alpha(u, v) - \beta| \leqslant \frac{\beta}{2t}$ for some $\beta > 0$, then (Y, α) is isometrically embeddable into l_2.*

We add a vector O (so that we have $k(T+1)+1$ vectors) and set the following distance requirements.

1. $\|v(i)_1 - O\|_2 = \sqrt{\mu}$ for all u_i.
2. $\|u(i)_1 - v(i')_1\|_2 = \sqrt{\mu - 2\rho(u(i), v(i')) + 2\nu_S(u(i), v(i'))}$ for all $u(i), v(i')$.

Note that $|\rho(u(i), v(i')) - \nu_S(u(i), v(i'))| \leqslant \frac{(1-\mu)^L}{T+1} \leqslant \frac{\mu}{5(T+1)k}$, where the last inequality follows from Theorem 5.2. This implies

$$|\|u(i)_1 - v(i')_1\|_2 - \sqrt{\mu}| \leqslant \sqrt{\mu}(1 - \sqrt{1 - \frac{1}{2.5(T+1)k}}) \leqslant \sqrt{\mu} \cdot \frac{1}{2((T+1)k+1)}.$$

By Lemma B.8, there are vectors $\{u(i)_1, v(i)_1\}_i$ and O that meet the above distance requirements. Without loss of generality, assume that O is the origin. Defining $u(i) := u(i)_1 \oplus u(i)_2$ satisfies

1. $\|u(i)\|_2^2 = \mu + \frac{1}{T+1}$.
2. $u(i) \cdot v(i') = u(i)_1 \cdot v(i')_1 + u(i)_2 \cdot v(i')_2 = \frac{2\mu - \|u(i)_1 - v(i')_1\|_2^2}{2} + \nu_S(u(i), v(i')) = \frac{\mu}{2} + \rho(u(i), v(i'))$.

□

B.4 Rounding

Lemma B.9 (Restatement of Lemma 5.8)**.** *There exists μ depending on T and ϵ such that, in the above rounding scheme, the probability that $l_V(u) = 0$ and $l_V(v) = l_A(u, v)$ is at most $\frac{1-12\epsilon}{T+1}$.*

Proof. For notational simplicity, assume $l_A(u, v) = 0$ — which is not allowed in actual instances. Then $u(i)$ and $v(i)$ become *matching* vectors — $\rho(u(i), v(i)) = \frac{1-\mu}{T+1} + \frac{\mu}{(T+1)^2}$ and $\rho(u(i), v(j)) = \frac{\mu}{(T+1)^2}$ for $i \neq j$. The following is the list of all possible inner products between $2(T+1)$ vectors.

1. $\|u(i)\|_2^2 = \mu + \frac{1}{T+1}$.
2. $u(i) \cdot u(j) = \frac{\mu}{2}$ for $i \neq j$.
3. $u(i) \cdot v(i) = \frac{\mu}{2} + \frac{1-\mu}{T+1} + \frac{\mu}{(T+1)^2}$.
4. $u(i) \cdot v(j) = \frac{\mu}{2} + \frac{\mu}{(T+1)^2}$ for $i \neq j$.

Even though we used Lemma B.8 as a black-box to obtain the current embedding, we can explicitly represent $u(i), v(i)$'s in the Euclidean space. They can be represented as a linear combination of $(T+1) + (T+1)^2 + 2(T+1) + 1$ orthogonal vectors (with different lengths), which can be classified into the following four categories:

- $a(i)$ for $i \in [T]$: Length $\sqrt{\frac{1-\mu}{T+1}}$. Denotes the event that (u, v) is not deleted and $l_V(u) = l_V(v) = i$.
- $b(i, j)$ for $i, j \in [T]$: Length $\sqrt{\frac{\mu}{(T+1)^2}}$. Denotes the event that (u, v) is deleted and $l_V(u) = i$, $l_V(v) = j$.
- $c(i), c'(i)$ for $i \in [T]$: Length $\sqrt{\frac{\mu}{2}}$. One of them is assigned for each of $2(T+1)$ vectors.

- d: Length $\sqrt{\frac{\mu}{2}}$. Common for all vectors.

Let

$$u(i) := a(i) + \sum_j b(i,j) + c(i) + d$$
$$v(i) := a(i) + \sum_j b(j,i) + c'(i) + d.$$

It is straightforward to check that the following representation of $u(i)$ and $v(i)$ satisfy all the inner product requirements.

For each vector $u(i)$, we denote the random variable equal to the inner product of $u(i)$ and g by $U(i)$. Similarly, define $V(i), A(i), B(i,j), C(i), C'(i), D(i)$ for $v(i), a(i), b(i,j), c(i), c'(i), d(i)$ respectively. Each random variable follows the Gaussian distribution with mean 0 and standard deviation same with the length of the corresponding vector. Furthermore, the inner products of two vectors is the same with the covariance of corresponding random variables. The following lemma shows that our consistent local distributions actually satisfy each edge with probability close to $\frac{1}{T+1}$, proving Theorem 1.4.

Lemma B.10. *Fix $i \in [T]$ and $0 < \epsilon < 1/24$. If $\mu \leqslant \frac{\epsilon^2}{256(T+1)\log^2(\frac{T+1}{\epsilon})}$,*

$$\Pr[l_V(u) = i, l_V(v) = i] \geqslant \frac{1 - 12\epsilon}{T+1}.$$

Proof. We compute the probability that u and v are assigned the same label i.

$$\begin{aligned}
& Pr[l_V(u) = i, l_V(v) = i] \\
\geqslant\ & \Pr[A(i) = \max_j[A(j)]] * \Pr[\max_j[\sum_k B(j,k)], \max_j[\sum_k B(k,j)], \max_j[C(j)], \max_j[C'(j)] \\
& \leqslant \frac{A(i) - \max_{j \neq i}[A(j)]}{4} | A(i) = \max_j[A(j)]] \\
\geqslant\ & \frac{1}{T+1} \Pr[\max_j[\sum_k B(j,k)], \max_j[\sum_k B(k,j)], \max_j[C(j)], \max_j[C'(j)] \leqslant \frac{\max_j[A(j)] - \max 2_j[A(j)]}{4}]
\end{aligned}$$

We argue that the above quantity is close to $\frac{1}{T+1}$ by showing that each of 4 quantities

$$\max_j[\sum_k B(j,k)], \max_j[\sum_k B(k,j)], \max_j[C(j)], \max_j[C'(j)]$$

is greater than $\frac{\max_j[A(j)] - \max 2_j[A(j)]}{4}$ with small probability. Note that $\sum_k B(j,k)$ follows the Gaussian distribution with mean 0 and variance $\frac{\mu}{T+1}$, which is much less than that of $C(j)$. Since $C(j)$ and $C'(j)$ follow the same distribution, it is enough to show that $\max_j[C(j)] > \frac{\max_j[A(j)] - \max 2_j[A(j)]}{4}$ with small probability. The following claim proves the lemma. □

Claim B.11. *Let $0 < \epsilon < 1/4$. If $\mu \leqslant \frac{\epsilon^2}{256(T+1)\log^2(\frac{T+1}{\epsilon})}$,*

$$\Pr[\max_j[C(j)] > \frac{\max_j[A(j)] - \max 2_j[A(j)]}{4}] < 3\epsilon.$$

Proof. The above probability can be rewritten as

$$\Pr[\sqrt{\frac{\mu}{2}}\max_j[g_j] > \sqrt{\frac{1-\mu}{T+1}}\frac{\max_j[g'_j]-\max 2_j[g'_j]}{4}]$$

where $g_0,\ldots,g_T,g'_0,\ldots,g'_T$ are independent standard Gaussian random variables.

Let $x \geqslant \sqrt{\mu \log \frac{T+1}{\epsilon}}$. By Lemma B.12,

$$\Pr[\sqrt{\frac{\mu}{2}}\max_j[g_j] > x] < \epsilon.$$

Let $x \leqslant \frac{\epsilon}{8\sqrt{\log\frac{T+1}{\epsilon}}}\sqrt{\frac{1-\mu}{T+1}}$. By Lemma B.13,

$$\Pr[\sqrt{\frac{1-\mu}{T+1}}\frac{\max_j[g'_j]-\max 2_j[g'_j]}{4} < x] < 2\epsilon.$$

The fact that $\mu \leqslant \frac{\epsilon^2}{256(T+1)\log^2(\frac{T+1}{\epsilon})}$ ensures that there is x that satisfies the both Lemma B.12 and B.13. Taking union bound proves the lemma. □

It remains to prove the following two lemmas about Gaussians. We prove them in Appendix E using some basic properties of Gaussians.

Lemma B.12. *Let $g_1,\ldots,g_n$ ($n \geqslant 2$) be independent standard Gaussian random variables and $0<\epsilon<1$. If $x \geqslant \sqrt{2\log\frac{n}{\epsilon}}$,*

$$\Pr[\max_j[g_j] \leqslant x] \geqslant 1-\epsilon.$$

Lemma B.13. *Let $g_1,\ldots,g_n$ ($n \geqslant 2$) be independent standard Gaussian random variables and $0<\epsilon<1/4$. If $x \leqslant \frac{\epsilon}{2\sqrt{\log\frac{n}{\epsilon}}}$,*

$$\Pr[\max_j[g_j]-\max_j 2[g_j] \geqslant x] \geqslant (1-2\epsilon).$$

□

C $(\frac{1}{4}+\Omega(\frac{1}{T}))$-Approximation Algorithm for Generalized Max-Dicut

In this section, we propose an approximation algorithm for Generalized Max-Dicut(T) that guarantees $(\frac{1}{4}+\frac{1}{16T})$ fraction of the optimal solution, proving Theorem 4.1. It is based on the 2-rounds of the Sherali-Adams hierarchy (also known as the standard LP), defined as the following:

$$\begin{array}{lll} \text{maximize} & \sum_{(u,v)\in A} x_{(u,v)}(0,l_A(u,v)) & \\ \text{subject to} & \sum_{\alpha\in[T]^S} x_S(\alpha)=1 & \text{for all } S\subseteq V, |S|\leqslant 2 \\ & \sum_{j\in[T]} x_{(u,v)}(i,j)=x_u(i) & \text{for all } u\neq v, i\in[T] \end{array}$$

The algorithm is almost identical to the simple $\frac{1}{4}$-approximation algorithm. For each vertex v, independently set $l_V(v)=0$ with probability $\frac{1+x_v(0)}{2}$, and $l_V(v)=i$ ($i\neq 0$) with probability $\frac{x_v(i)}{2}$. Equivalently, we assign each vertex 0 with probability half and follow its marginal x_v with probability half.

For each edge $(u,v) \in A$, let $c = c(u,v) := x_{(u,v)}(0, l_A(u,v))$ so that the value the solution $\{x_S(\alpha)\}$ to the LP is $\mathbb{E}_{(u,v)}[c(u,v)] \geqslant \mathsf{Opt}$. The probability that (u,v) is satisfied is

$$(\frac{1+x_u(0)}{2})(\frac{x_v(l_A(u,v))}{2}) \geqslant \frac{c}{4} + \frac{c^2}{4}$$

since $x_u(0), x_v(l_A(u,v)) \geqslant c$. Therefore, the expected fraction of satisfied edges is at least

$$\mathop{\mathbb{E}}_{(u,v)\in A}[\frac{c(u,v)}{4} + \frac{c(u,v)^2}{4}] \geqslant \frac{\mathsf{Opt}}{4} + \frac{\mathsf{Opt}^2}{4} \geqslant \frac{\mathsf{Opt}}{4} + \frac{\mathsf{Opt}}{16T}$$

since $\mathsf{Opt} \geqslant \frac{1}{4T}$ (focusing on the label with the most edges and finding the maximum dicut with respect to the edges with this label guarantees to satisfy $\frac{1}{4T}$ fraction of edges).

D Details in the Unique Games-Hardness

D.1 Lemmas about the Dictatorship Test

Lemma D.1 (Restatement of Lemma 4.4)**.** *For any t, $\rho(\Omega_1, \Omega_2; \mathbf{P}^t) \leqslant \sqrt{\frac{2}{T\delta}}$.*

Proof. Let $f : \Omega_1 \to \mathbb{R}$ be the function satisfying $\mathbb{E}[f] = 0$, $\mathbb{E}[f^2] = 1$. Let L be the Markov operator defined in Section 2.1 of Mossel [28] such that

$$(Lf)(y) = \mathbb{E}[f(X)|Y = y]$$

for $y \in \Omega_2$ and $(X,Y) \in \Omega_1 \times \Omega_2$ is distributed according to $\mathbf{P}^t$. By Lemma 2.8 of [28],

$$\rho(\Omega_1, \Omega_2) = \sup_f \sqrt{\mathbb{E}[(Lf)^2]}.$$

Let $f(i) = a_i, (Lf)(i) = b_i$ for $i \in [T]$. $b_t = a_0$ and all the other b_i's are equal to $\mathbb{E}_{\mathbf{P}_1'}[f]$, which is equal to $(\frac{1}{1-\frac{1-\delta}{T}})(\mathbb{E}_{\mathbf{P}_1}[f] - \frac{1-\delta}{T}a_0) = (\frac{1}{1-\frac{1-\delta}{T}})(-\frac{1-\delta}{T}a_0)$.

$$\begin{aligned}
\mathbb{E}[(Lf)^2] &= (\frac{1-\delta}{T})a_0^2 + (1-\frac{1-\delta}{T})[(\frac{1}{1-\frac{1-\delta}{T}})(-\frac{1-\delta}{T}a_0)]^2 \\
&= (\frac{1-\delta}{T})a_0^2 + (\frac{1}{1-\frac{1-\delta}{T}})(\frac{1-\delta}{T}a_0)^2 \\
&= (\frac{1-\delta}{T})a_0^2[1 + (\frac{1}{1-\frac{1-\delta}{T}})(\frac{1-\delta}{T})] \\
&\leqslant \frac{2}{T}a_0^2 \\
&\leqslant \frac{2}{T\delta}
\end{aligned}$$

Since $\delta a_0^2 \leqslant \mathbb{E}[f^2] \leqslant 1$. ∎

D.2 Reduction From Unique Games

In this subsection, we introduce the reduction from the Unique Games to Generalized Max-Dicut(T), using the dictatorship test constructed. We first introduce the Unique Games Conjecture [20], which is stated below.

Definition D.2 (Unique Games)**.** *An instance $\mathcal{L}(G(U \cup W, E), [R]^+, \{\pi(u,w)\}_{(u,w)\in E})$ of Unique Games consists of a regular bipartite graph $G(U \cup W, E)$ and a set $[R]^+$ of labels. For each edge $(u,w) \in E$ there is a constraint specified by a permutation $\pi(u,w) : [R]^+ \to [R]^+$. Given a labeling $l : U \cup W \to [R]^+$, let $\mathsf{Val}_{\mathsf{UG}}(l)$ be the fraction of labels satisfied by l, where an edge $e = (u,w)$ is said to be satisfied if $l(u) = \pi(u,w)(l(w))$. Let $\mathsf{Opt}_{\mathsf{UG}}(\mathcal{L}) = \max_l(\mathsf{Val}_{\mathsf{UG}}(l))$.*

Conjecture D.3 (Unique Games Conjecture [20])**.** *For any constant $\alpha > 0$, there is $R = R(\alpha)$ such that, for a Unique Games instance $\mathcal{L}$ with label set $[R]^+$, it is NP-hard to distinguish between*

- $\mathsf{Opt}_{\mathsf{UG}}(\mathcal{L}) \geqslant 1 - \alpha$.
- $\mathsf{Opt}_{\mathsf{UG}}(\mathcal{L}) \leqslant \alpha$.

Theorem D.4 (Restatement of Theorem 1.3)**.** *Under the Unique Games Conjecture, it is NP-hard to approximate Generalized Max-Dicut(T) on directed acyclic graphs within a factor of $\frac{1}{4} + O(\frac{1}{T^{1/4}})$.*

Proof. Given an instance of $\mathcal{L}(G(U \cup W, E), [R]^+, \{\pi(v,w)\}_{(v,w)\in E})$ of Unique Games, we construct an instance $\mathcal{D}(\mathcal{V}, \mathcal{A}), l_{\mathcal{A}}$ of Generalized Max-Dicut(T). For $x \in [T]^R$ and a permutation $\pi : [R]^+ \to [R]^+$, let $x \circ \pi \in [T]^R$ be defined by $(x \circ \pi)_i = (x)_{\pi^{-1}(i)}$. Let $D = (V, A)$ be the fixed-size graph where the maximum dicut has at most $(\frac{1}{4} + \frac{1}{T^{1/4}})$ fraction of edges.

- $\mathcal{V} = U \times V \times [T]^R$.
- Sample $w \in W$ uniformly at random and its neighbors u_1, u_2 uniformly and independently. Sample $t \in [T]^+$, $(v_1, v_2) \in A$, and $x, y \in [T]^R$ from the dictatorship test. Add an edge $((u_1, v_1, x \circ \pi_{u_1,w}), (u_2, v_2, y \circ \pi_{u_2,w}))$ to $\mathcal{A}$ with label t. The weight is equal to the probability that this edge is sampled.

Completeness Suppose that $\mathsf{Val}_{\mathsf{UG}}(l) \geqslant 1 - \alpha$ for some labeling $l : U \cup W \to [R]^+$.

Set $l_{\mathcal{V}}(u, v, (x_1, \ldots, x_R)) = x_{l(u)}$. For w, u_1, u_2 sampled as above, with probability $1-2\alpha$, $\pi(u_1,w)^{-1}(l(u_1)) = \pi(u_2,w)^{-1}(l(u_2))$. In that case, by Lemma 4.5,

$$\begin{aligned}
&\Pr_{v_1,v_2,t,x,y}[l_{\mathcal{V}}(u_1, v_1, x \circ \pi_{u_1,w}) = 0, l_{\mathcal{V}}(u_2, v_2, y \circ \pi_{u_2,w}) = t] \\
=&\Pr_{v_1,v_2,t,x,y}[(x \circ \pi_{u_1,w})_{l(u_1)} = 0, (y \circ \pi_{u_2,w})_{l(u_2)} = t] \\
=&\Pr_{v_1,v_2,t,x,y}[(x)_{\pi(u_1,w)^{-1}(l(u_1))} = 0, (y)_{\pi(u_2,w)^{-1}(l(u_2))} = t] \\
\geqslant&\frac{1-\delta}{T}.
\end{aligned}$$

Therefore, $\mathsf{Val}_{\mathsf{GMD}}(l_{\mathcal{V}}) \geqslant \frac{(1-2\alpha)(1-\delta)}{T}$.

Soundness For each $u \in U, v \in V$ and $t \in [T]$, let $F_{u,v,t} : [T]^R \to \{0,1\}$ be defined by

$$F_{u,v,t}(x) = 1 \text{ if and only if } l_{\mathcal{V}}(u,v,x) = t.$$

Similarly, for each $w \in W, v \in V$ and $t \in [T]$, let $H_{w,v,t} : [T]^R \to [0,1]$ be the function defined by

$$H_{w,v,t}(x) = \mathop{\mathbb{E}}_{(u,w)\in E}[F_{u,v,t}(u, x \circ \pi(u,w))] = \Pr_{(u,w)\in E}[l_{\mathcal{V}}(u,v,x \circ \pi(u,w)) = t].$$

Suppose that there exists l_V such that $\mathsf{Val}_{\mathsf{GMD}}(l_V) \geqslant \frac{1}{4T} + \frac{5}{T^{5/4}}$. For at least $\frac{1}{T^{5/4}}$ fraction of w, an edge of $\mathcal{A}$ sampled by first choosing w is satisfied with probability more than $\frac{1}{4T} + \frac{4}{T^{5/4}}$. By Lemma 4.6, there exist τ and D, such that, for each such w, we have $\mathsf{Inf}_i^{\leqslant d}[H_{w,v,t}] > \tau$ for some i, v and t. Set $l_{\mathcal{V}}(w) = i$. For other w's, choose $l_V(w)$ arbitrarily.

From the representation of influences in terms of Fourier coefficients (see Khot et al. [21]),

$$\tau < \mathsf{Inf}_i^{\leqslant d}[H_{w,v,t}] \leqslant \mathop{\mathbb{E}}_{(u,w)\in E}[\mathsf{Inf}^{\leqslant d}_{\pi_{u,w}(i)}[F_{u,v,t}]]$$

and we conclude that $\tau/2$ fraction of neighbors u of w have $\mathsf{Inf}^{\leqslant d}_{\pi_{u,w}(i)}(F_{u,v,t}) \geqslant \tau/2$. We choose $l_{\mathcal{V}}(u)$ uniformly from

$$\left\{ i : \mathsf{Inf}_i^{\leqslant d}[F_{u,v,t}] \geqslant \tau/2 \text{ for some } t, v \right\}.$$

Since $\sum_i \mathsf{Inf}_i^{\leqslant D}[F_{u,t}] \leqslant d$, there are at most $\frac{2(T+1)d|V|}{\tau}$ of candidate i's for each u. If u have no candidate, choose $l_V(u)$ arbitrarily. The above strategy satisfies $(\frac{1}{T^{5/4}})(\frac{\tau}{2})(\frac{\tau}{2(T+1)D|V|})$ fraction of constraints in expectation. Taking α small enough completes the proof of the theorem. □

Now, we present the full proof of our main theorem.

Theorem D.5 (Restatement of Theorem 1.1)**.** *Under the Unique Games Conjecture, for any $\epsilon > 0$, it is NP-hard to approximate Graph Pricing within a factor of $\frac{1}{4} + \epsilon$.*

Proof. Given $\epsilon > 0$, let T large enough so that $\frac{1}{T^{1/4}} < \frac{\epsilon}{2}$. Theorem 1.3 tells that it is hard to distinguish

- Completeness: $\mathsf{Opt}_{\mathsf{GMD}} \geqslant \frac{1}{T} - \frac{2}{T^{4/5}} = \frac{1-O(\epsilon)}{T}$.
- Soundness: $\mathsf{Opt}_{\mathsf{GMD}} \leqslant \frac{1}{4T} + \frac{5}{T^{4/5}} = \frac{1+O(\epsilon)}{4T}$.

Let $t = \frac{T}{\epsilon}$. We can assume that each vertex in the Unique Games instance is of degree at least t, since duplicating each vertex v into t copies $v_1, ..., v_t$ and duplicating each constraint (u, v) into t^2 copies $(u_i, v_j)_{1\leqslant i,j\leqslant t}$ preserves the optimum. Therefore, the instance of Generalized Max-Dicut obtained from the above Unique Games instance will have $\mathsf{ndeg} \geqslant t$. Theorem 3.1 shows that it is NP-hard to distinguish

- Completeness: $\mathsf{Opt}_{\mathsf{GP}} \geqslant \mathsf{Opt}_{\mathsf{GMD}} = \frac{1-O(\epsilon)}{T}$.
- Soundness: $\mathsf{Opt}_{\mathsf{GP}} \leqslant \mathsf{Opt}_{\mathsf{GMD}} + \frac{1}{t} = \frac{1+O(\epsilon)}{4T} + \frac{\epsilon}{T} = \frac{1+O(\epsilon)}{4T}$.

□

E Proofs of Lemmas about Gaussians

Let $\phi(x)$ and $\Phi(x)$ be the probability density function (PDF) and the cumulative distribution function (CDF) of the standard Gaussian, respectively. Let $\tilde{\Phi}(x) = 1 - \Phi(x)$. We begin with the following simple fact about the tail of Φ.

Lemma E.1 ([8])**.** *For any $t > 0$, $\frac{t}{\sqrt{2\pi}(t^2+1)} e^{-\frac{t^2}{2}} < \tilde{\Phi}(t) < \frac{1}{\sqrt{2\pi}t} e^{-\frac{t^2}{2}}$.*

Lemma E.2 (Restatement of Lemma B.12)**.** *Let $g_1, \ldots, g_n$ ($n \geqslant 2$) be independent standard Gaussian random variables and $0 < \epsilon < 1$. If $x \geqslant \sqrt{2 \log \frac{n}{\epsilon}}$,*

$$\Pr[\max_j [g_j] \leqslant x] \geqslant 1 - \epsilon.$$

Proof. Note that $x \geqslant \sqrt{2 \log 2}$, so $\frac{1}{\sqrt{2\pi}x} \leqslant 1$.

$$\begin{aligned} & x \geqslant \sqrt{2 \log \frac{n}{\epsilon}} \\ \Rightarrow \quad & \frac{1}{\sqrt{2\pi}x} \exp(-\frac{x^2}{2}) \leqslant \frac{\epsilon}{n} \\ \Rightarrow \quad & 1 - \Phi(x) \leqslant \frac{\epsilon}{n}, \end{aligned}$$

where the last inequality follows from Lemma E.1. We can conclude that

$$\Pr[\max_j [C(j)] \leqslant x] = \Phi(x)^n \geqslant (1 - \frac{\epsilon}{n})^n \geqslant 1 - \epsilon.$$

□

Lemma E.3 (Restatement of Lemma B.13)**.** *Let $g_1, \ldots, g_n$ ($n \geqslant 2$) be independent standard Gaussian random variables and $0 < \epsilon < 1/4$. If $x \leqslant \frac{\epsilon}{2\sqrt{\log \frac{n}{\epsilon}}}$,*

$$\Pr[\max_j [g_j] - \max_j 2[g_j] \geqslant x] \geqslant (1 - 2\epsilon).$$

Proof.

$$\begin{aligned} \Pr[\max_j [g_j] - \max_j 2[g_j] \geqslant x] \quad &\geqslant \quad n \int_{-\infty}^{\infty} \Phi[y - x]^{n-1} \phi(y) dy \\ &\geqslant \quad n \int_{-\infty}^{b} \Phi[y - x]^{n-1} \phi(y) dy \quad \text{for some } b \text{ fixed later} \\ &= \quad n \int_{-\infty}^{b} \Phi[y - x]^{n-1} \phi(y - x) \frac{\phi(y)}{\phi(y - x)} dy \\ &\geqslant \quad (\inf_{y \in [-\infty, b]} \frac{\phi(y)}{\phi(y - x)}) \int_{-\infty}^{b} n \Phi[y - x]^{n-1} \phi(y - x) dy \\ &= \quad (\inf_{y \in [-\infty, b]} \frac{\phi(y)}{\phi(y - x)}) \int_{-\infty}^{b} (\Phi[y - x]^n)' dy \\ &= \quad (\inf_{y \in [-\infty, b]} \frac{\phi(y)}{\phi(y - x)}) \Phi[b - x]^n \end{aligned}$$

Let $b = x + \sqrt{2\log\frac{n}{\epsilon}}$. By the same argument with Lemma B.12, we have

$$\begin{aligned} & 1 - \Phi[b-x] \leqslant \frac{\epsilon}{n} \\ \Rightarrow \quad & \Phi[b-x] \geqslant 1 - \frac{\epsilon}{n} \\ \Rightarrow \quad & \Phi[b-x] \geqslant (1-\epsilon)^{1/n} \end{aligned}$$

Now we bound

$$\inf_{y\in[-\infty,b]} \frac{\phi(y)}{\phi(y-x)} = \inf_{y\in[-\infty,b]} \exp(-\frac{y^2}{2} + \frac{(y-x)^2}{2}) = \inf_{y\in[-\infty,b]} \exp(\frac{-2xy+x^2}{2}) = \exp(\frac{-2bx+x^2}{2})$$

where the last inequality holds since it is monotonically decreasing in y. $x \leqslant \frac{\epsilon}{2\sqrt{\log\frac{n}{\epsilon}}}$ implies

$$\begin{aligned} & x(x + \sqrt{2\log\frac{n}{\epsilon}}) \leqslant \epsilon \\ \Rightarrow \quad & bx \leqslant \epsilon \\ \Rightarrow \quad & \frac{-2bx+x^2}{2} \geqslant -\epsilon \\ \Rightarrow \quad & \exp(\frac{-2bx+x^2}{2}) \geqslant \exp(-\epsilon) \geqslant 1-\epsilon \end{aligned}$$

Since both $\inf_{y\in[-\infty,b]} \frac{\phi(y)}{\phi(y-x)}$ and $\Phi[b-x]^n$ are at least $1-\epsilon$, the lemma follows. □

Lemma E.4 (Restatement of Lemma 4.8)**.** *Fix $\rho, \alpha \in (0,1)$. The function $f(x) := \Gamma_\rho(\alpha, x)$ is concave.*

Proof. Let Y, Z be independent Gaussians and $X := \rho Y + \sqrt{1-\rho^2}Z$. Fix $0 \leqslant a \leqslant b$. We will show that $f(a) + f(b) \geqslant f(a+b)$. Let $x = \tilde{\Phi}^{-1}(a+b), y = \tilde{\Phi}^{-1}(b), z = \tilde{\Phi}^{-1}(a), w = \tilde{\Phi}^{-1}(\alpha)$. Note that $x \leqslant y \leqslant z$.

$$\begin{aligned} & f(a) + f(b) - f(a+b) \\ =& \Pr[Y \geqslant y \text{ and } X \geqslant w] + \Pr[Y \geqslant z \text{ and } X \geqslant w] - \Pr[Y \geqslant x \text{ and } X \geqslant w] \\ =& \Pr[Y \geqslant z \text{ and } X \geqslant w] - \Pr[x \leqslant Y \leqslant y \text{ and } X \geqslant w] \\ \geqslant& \Pr[Y \geqslant z \text{ and } Z \geqslant \frac{w-\rho z}{\sqrt{1-\rho^2}}] - \Pr[x \leqslant Y \leqslant y \text{ and } Z \geqslant \frac{w-\rho y}{\sqrt{1-\rho^2}}] \\ =& a(\Pr[Z \geqslant \frac{w-\rho z}{\sqrt{1-\rho^2}}] - \Pr[Z \geqslant \frac{w-\rho y}{\sqrt{1-\rho^2}}]) \\ \geqslant& 0 \end{aligned}$$

□

Lemma E.5 (Restatement of Lemma 4.9)**.** *For large enough T and $\delta = \frac{1}{T^{1/4}}$, the following holds. For any $a \in [0,1], b \in [0, \frac{1}{T}]$ and $\rho \in (0, \sqrt{\frac{2}{T\delta}})$, $\Gamma_\rho(a,b) \leqslant ab + \frac{2}{T^{5/4}}$.*

Proof. Let Y, Z be independent Gaussians and $X := \rho Y + \sqrt{1-\rho^2}Z$. Let $x = \tilde{\Phi}^{-1}(a)$ and $y = \tilde{\Phi}^{-1}(b)$. By taking $T > 2$, we can assume $b < \frac{1}{2}$ and $y > 0$, while we do not put any assumption on a and x.

$$\begin{aligned}\Gamma_\rho(a,b) &= \Pr[X \geqslant x \text{ and } Y \geqslant y] \\ &\leqslant \Pr[Z \geqslant \frac{x-2\rho y}{\sqrt{1-\rho^2}} \text{ and } y \leqslant Y \leqslant 2y] + \Pr[Y \geqslant 2y] \\ &\leqslant \Pr[Z \geqslant \frac{x-2\rho y}{\sqrt{1-\rho^2}} \text{ and } Y \geqslant y] + \Pr[Y \geqslant 2y] \\ &\leqslant b \cdot \tilde{\Phi}(\frac{x-2\rho y}{\sqrt{1-\rho^2}}) + \tilde{\Phi}(2y). \end{aligned} \tag{1}$$

By Lemma E.1, $\tilde{\Phi}(2y) < \frac{1}{2\sqrt{2\pi}y}\exp(-2y^2) < b^3 < \frac{1}{T^{5/4}}$.

- $a \geqslant 1 - \frac{1}{T^{1/4}}$: (1) is bounded by $b + \frac{1}{T^{5/4}} \leqslant (a + \frac{1}{T^{1/4}})b + \frac{1}{T^{5/4}} \leqslant ab + \frac{2}{T^{5/4}}$.
- $b \leqslant \frac{1}{T^{5/4}}$: (1) is bounded by $b + \frac{1}{T^{5/4}} \leqslant \frac{2}{T^{5/4}}$.
- $a \leqslant 1 - \frac{1}{T^{1/4}}$ and $b \geqslant \frac{1}{T^{5/4}}$: Note that $x \geqslant -10\sqrt{\log T}$ and $y \leqslant 10\sqrt{\log T}$. Since $\rho \leqslant \sqrt{\frac{2}{T\delta}} = \frac{\sqrt{2}}{T^{3/8}}$,

  $$(x-2\rho y) - \sqrt{1-\rho^2}(x - \frac{1}{T^{1/4}}) \geqslant \begin{cases} -2\rho y + \frac{1}{2T^{1/4}} \geqslant 0 & \text{if } x \geqslant 0 \\ \rho^2 x - 2\rho y + \frac{1}{2T^{1/4}} \geqslant 0 & \text{if } -10\sqrt{\log T} \leqslant x \leqslant 0, \end{cases}$$

  which shows that $\frac{x-2\rho y}{\sqrt{1-\rho^2}} \geqslant x - \frac{1}{T^{1/4}}$. Therefore,

  $$(1) \leqslant b \cdot \tilde{\Phi}(x - \frac{1}{T^{1/4}}) + \frac{1}{T^{5/4}} \leqslant b(a + \frac{1}{T^{1/4}}) + \frac{1}{T^{5/4}} \leqslant ab + \frac{2}{T^{5/4}},$$

  where the second inequality follows from $\phi(x) \leqslant 1$ for all $x \in \mathbb{R}$.

□